18.721: Introduction to Algebraic Geometry

Lecturer: Professor Mike Artin

Notes by: Andrew Lin

Spring 2020

1 February 3, 2020

Algebraic geometry is a beautiful subject, and it's usually taught as a mid-level graduate course, so we'll need to discuss things in this class without a lot of background. In particular, we won't assume commutative algebra (18.705), though that might be useful. (18.702 is essential, though.)

Throughout this class, we'll work with the scalars $\mathbb{C}$ – algebraic geometry can be done with any field of scalars, but we're making this choice to give a bit of intuition into the geometry. (Real numbers don't work, by the way – we need algebraic closure.)

The central objects of study here are **systems of polynomial equations** in multiple variables $x = (x_1, x_2, \cdots, x_n)$ of the form $\{f_1(x) = 0, \cdots, f_k(x) = 0\}$. **Geometrically**, this means we'll look at the locus of solutions X in **affine space** $\mathbb{A}^n = \mathbb{C}^n$, and **algebraically**, it means we'll look at the polynomial ring $\mathbb{C}[x]$ (that is, polynomials in our variables x_i), modded out by our polynomials, which yields an **algebra** of the form

$$A = \mathbb{C}[x]/(f_1, \cdots, f_k).$$

Note that any ring that can be generated by finitely many elements will be of this form – being finitely generated means that there's a surjective map $\mathbb{C}[x] \to A$, and the Hilbert Basis Theorem tells us that the kernel can always be generated by finitely many elements.

Remember that the **ideal** of $\mathbb{C}[x]$ generated by $f_1, \cdots, f_k$ is the set of linear combinations $g_1 f_1 + \cdots + g_k f_k$, where $g_i \in \mathbb{C}[x]$. The point is that the geometry and algebra of this situation will complement each other!

We'll start with some simple examples:

Example 1

Consider the two-variable polynomial $f(x, y) = y^2 - x^3 - x^2$.

Let X be the locus (of zeros), described via $y^2 = x^3 + x^2$. X lives in 4-space (because x and y can be complex), but we can at least graph it in real space:

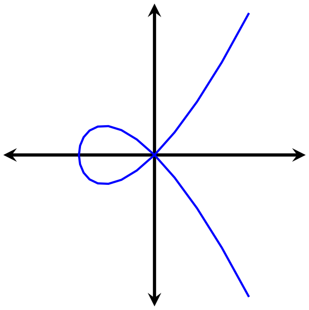

The geometry is reflected in the algebra here, because we can actually parameterize this curve using polynomials: if we draw a line of slope t from the origin (which is a double zero), then $y = tx$, so

$$f(x, tx) = t^2x^2 - x^3 - x^2 = x^2(t^2 - x - 1),$$

so $x = t^2 - 1$, which tells us that $y = t^3 - t$. In algebraic terms, this means that we can construct a map $\mathbb{A}^1_t \to X$ (the one-dimensional affine space maps to our locus X). In addition, we can take the algebra $A = \mathbb{C}[x, y]/(f)$ and map it to $\mathbb{C}[t]$ (though these are not isomorphic algebras). Really, what's going on is that we take the one-dimensional complex t and identifying two points together, and that gives us our zero locus.

Example 2

On the other hand, consider the curve $f(x, y) = y^2 - x^3 + x$.

Now we want $y^2 = x^3 - x$, but notice that the (real) curve that comes out of this is no longer just one piece – things are more complicated here. We'll come back to this in a second.

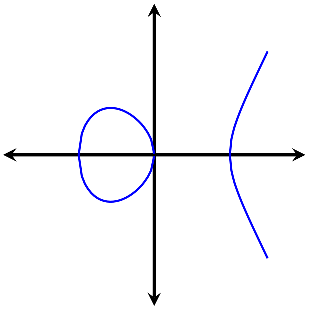

Affine space is interesting, but it's often important to study **projective space**:

Definition 3

Points in **projective space** $\mathbb{P}^n$ are classes of nonzero vectors $(x_0; x_1, \cdots, x_n)$ with the equivalence relation that

$$(x_0; x_1, \cdots, x_n) = (\lambda x_0; \cdots, \lambda x_n)$$

for any $\lambda \in \mathbb{C}, \lambda \neq 0$.

One reason this is important is that $\mathbb{P}^n$ **is compact**, while $\mathbb{A}^n$ is not! Let's verify this:

Proof. Let V be $\mathbb{C}^{n+1}$, so then $V - \{0\}$ maps to $\mathbb{P}^n$. Then equivalence classes of $\mathbb{P}^n$ map to 1-dimensional subspaces of V, so we can look at the unit sphere S in V. (Since V is a $2n+2$-dimensional real space, S has dimension $2n+1$.) S is the set of points x such that $\sum \overline{x}_i x_i = 1$; since we're modding out the equivalence relation, every point in projective space is represented by points in S (not uniquely). So there is a surjective map from S to $\mathbb{P}^n$. Clearly S is closed and bounded, so it is compact, and therefore $\mathbb{P}^n$ is compact as well. □

So how can we find loci in projective space $\mathbb{P}^n$? We want polynomials $f(x_0, \cdots, x_n)$ so that if $f(x) = 0$, then $f(\lambda x) = 0$ as well for all complex λ. If we write out

$$f(x) = f_0 + f_1 + f_2 + \cdots + f_d$$

(splitting up our polynomial by degree), then

$$f(\lambda x) = f_0 + \lambda f_1 + \lambda^2 f_2 + \cdots + \lambda^d f_d.$$

For any given x where $f(x) = 0$, we can substitute in the values of $f_0(x)$, $f_1(x)$, and so on into our equation. We now have a polynomial in λ that needs to be satisfied for all complex values of λ: this is only possible if the polynomial is the zero polynomial! So that means $f_i(x) = 0$ for all i.

In other words, in $\mathbb{P}^n$, we should only be studying the zeros of **homogeneous polynomial equations** in $x = x_0, \cdots, x_n$ – this tells us the same information as if we try work with polynomials in general.

So how can we study a polynomial like $f(x, y) = y^2 - x^3 + x$ in projective space if it's not homogeneous? The solution is to use the extra variable to homogenize into a cubic

$$f(x, y, z) = y^2 z - x^3 + xz^2.$$

Let's discuss the locus of this polynomial in $\mathbb{P}^2_{x,y,z}$. First, let's go back to a more generic description: any point (x_0, x_1, x_2) can be identified with $(1, u_1, u_2)$, where $u_i = \frac{x_i}{x_0}$, as long as $x_0 \neq 0$. This gives us a subset $U^0 \subset \mathbb{P}^2$ of points, corresponding to points (u_1, u_2) in the affine plane $\mathbb{A}^2_{u_1,u_2}$. Meanwhile, the points (x_0, x_1, x_2) where $x_0 = 0$ are on a line L^0 (defined by $\{x_0 = 0\}$). So $\mathbb{P}^2 = U^0 \cup L^0$: U^0 is often called the **points at finite distance**, and L^0 is often called the **line at infinity**.

With this, we can go back to our equation $y^2 z - x^3 + xz^2 = 0$: relabeling, we can say that

$$x_2^2 x_0 - x_1^3 + x_1 x_0^2 = 0.$$

If we consider $x_0 = 1$, we get the affine space curve $x_2^2 - x_1^3 + x_1$ (which is what we started with). But if $x_0 = 0$, we see what happens at infinity: that just gives us x_1^3, which means that there's only one point at infinity, $(0, 0, 1)$, on the curve. (This has to do with being able to **compactify** the locus in one point.)

So let's compute the Euler characteristic of this locus X (the number of vertices plus the number of faces, minus the number of edges). We have four **branch points** (the three zeros on the x_1-axis, plus the point at infinity), but at every other point there are two different values of x_2. This means that X covers most of $\mathbb{P}^1$ twice, so we should triangulate the two copies of $\mathbb{P}^1$ instead. This gives two edges, faces, and vertices over every edge, face, and vertex, except at the four branch points (vertices), which are double-counted:

$$e(X) = 2e(\mathbb{P}^1) - 4 = 2 \cdot 2 - 4 = 0.$$

So the **genus** of the surface here satisfies

$$e = 2 - 2g \implies g(X) = 1.$$

This means that if we drew the locus in complex space, we see a torus (except missing the point at infinity).

2 February 5, 2020

We'll start by taking another look at our picture of the (real) projective plane, which can be described with points in three dimensions. Let's take the plane containing the points $(1, 0, 0)$, $(0, 1, 0)$, and $(0, 0, 1)$, which has equation $x_0 + x_1 + x_2 = 1$. Every one-dimensional subspace of $\mathbb{R}^3$ goes through this plane exactly once, except the subspaces with $x_0 + x_1 + x_2 = 0$.

Yesterday, we drew the curve $x_0x_2^2 = x_1^3 - x_0^2x_1$: there are three points where $x_2 = 0$, which are at $x_1 = 0, \pm x_0$. In addition, the curve intersects the line at infinity $x_0 = 0$ with multiplicity 3 at the point $(0, 1, 0)$, so that yields a "flex point" at the intersection:

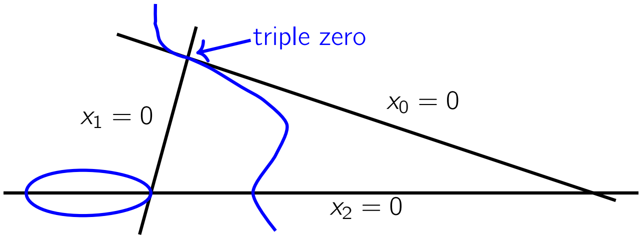

Remark 4. *Note that a point like $(-1, -1, 1)$ and a point like $(1, 1, -1)$ are the same in projective space, so the regions "wrap back" around.*

Another way to draw this real projective plane is to take a 2-sphere and identify opposite points with each other, but we won't use that picture here.

With that, let's move on to **projective plane curves**. We want to look at the locus of zeros for polynomials f, but remember here that we want to make sure f is an **irreducible homogeneous** polynomial.

The simplest case is a line: it takes the form

$$f = s_0x_0 + s_1x_1 + s_2x_2 = 0.$$

Another description is that we can take two points $p = (p_0, p_1, p_2)$ and $q = (q_0, q_1, q_2)$ and join them together to get a line $L_{pq} = \{ap + bq\}$, which corresponds to a one-dimensional subspace $(a, b) \in \mathbb{P}^1$. So a line looks like a one-dimensional projective space.

A **conic** is next: it's a combination of the degree-two monomials, which are $x_0^2, x_1^2, x_2^2, x_0x_1, x_0x_2, x_1x_2$. First, let's see how to do **linear** changes of coordinates in $\mathbb{P}^2$: if $x = (x_0, x_1, x_2)$, we just let $x' = Px$ for some 3×3 complex invertible matrix.

Proposition 5

For any conic C, we can choose coordinates so that the equation of C in new coordinates is $x_0x_1 + x_0x_2 + x_1x_2 = 0$.

So there's only one conic up to change in coordinates in projective space!

Proof. Choose three points on C, not on a line. (Being able to do this is the first problem on our homework, due next week.) Adjust coordinates so that these are $(1, 0, 0), (0, 1, 0)$, and $(0, 0, 1)$.

Then the modified function f satisfies $f(1, 0, 0) = 0$, so the coefficient of x_0^2 is 0. The same holds for x_1^2 and x_2^2, so now f is of the form $ax_0x_1 + bx_0x_2 + cx_1x_2$. And now scale the variables so that a, b, c become 1. □

Next, let's talk a bit about the tangent space: let C be a curve corresponding to the locus of zeros for f, and let $p \in C$ be a **smooth** point (this means the partial derivatives $\frac{\partial f}{\partial x_i}(p)$ are not all zero – the other case is called a **singular** point). For example, $x_0^3 + x_1^3 + x_2^3 = 0$ is smooth (except at the origin, which is not a point of projective space), but $x_0x_2^2 - x_1^3 - x_0x_1^2$ has partial derivatives $(x_2^2 - x_1^2, -3x_1^2 - 2x_0x_1, 2x_0x_2)$, which is zero for example at $(1, 0, 0)$.

There's a nice formula we can use:

Theorem 6 (Euler)

If f is homogeneous with degree d, then

$$x_0f_0 + x_1f_1 + x_2f_2 = d \cdot f.$$

where f_0 refers to the partial derivative of f with respect to x_0.

We can just check this for a monomial: for example, if $f = x_0^2x_1$, then

$$x_0f_0 + x_1f_1 + x_2f_2 = 2x_0^2x_1 + x_0^2x_1 = 3f.$$

(x_if_i contributes the degree of x_i towards the count of f.) So let's write down the tangent line of a curve C at a point p. To do that, pick another point q distinct from p, and define the line (this is equivalent to $\{ap + bq\}$ but better suited for our purposes)

$$L_{pq} = \{p + qt\} \cup \{q\}.$$

Let's now expand $f(p + qt)$ in a Taylor series:

$$f(p + qt) = f(p) + \left(\sum f_i(p)q_i\right) t + \frac{1}{2}\left(\sum_{i,j} q_if_{ij}q_j\right) t^2 + O(t^3).$$

Looking term by term now lets us know when $f(p + qt)$ is a tangent line. p is on the curve, so $f(p)$ is always zero. But the first derivative tells us about the slope, so we want to control the derivatives. We can write down the gradient

$$\nabla f = (f_0, f_1, f_2)$$

and the **Hessian matrix**

$$H = (f_{ij}), \quad f_{ij} = \frac{\partial^2 f}{\partial x_i \partial x_j}.$$

If we write our two points $p = \begin{bmatrix} p_0 \\ p_1 \\ p_2 \end{bmatrix}$ and $q = \begin{bmatrix} q_0 \\ q_1 \\ q_2 \end{bmatrix}$ as column vectors, we can now rewrite

$$L_{pq} = f(p) + \nabla_p qt + \frac{1}{2}(q^t H_p q)t^2 + O(t^3).$$

Define a bilinear form on column vectors in $\mathbb{C}^3$ via

$$\langle u, v \rangle = u^t H_p v,$$

and now we can rewrite all of our coefficients for our line: Euler's formula applied to the gradient tells us that

$$\nabla_p \cdot p = d \cdot f(p),$$

and Euler's formula applied to the Hessian tells us that

$$p^t H_p = \begin{bmatrix} p_0 & p_1 & p_2 \end{bmatrix} \begin{bmatrix} f_{00} & f_{01} & f_{02} \\ f_{10} & f_{11} & f_{12} \\ f_{20} & f_{21} & f_{22} \end{bmatrix} = (d-1) \begin{bmatrix} f_0 & f_1 & f_2 \end{bmatrix} = (d-1)\nabla_p.$$

Therefore,

$$p^t H_p p = (d-1)\nabla_p p = d(d-1)f(p),$$

and now L_{pq} can be written as

$$L_{pq} = \frac{1}{d(d-1)}\langle p, p\rangle + \frac{1}{d-1}\langle p, q\rangle t + \frac{1}{2}\langle q, q\rangle t^2 + O(t^3).$$

p is on our curve C if $\langle p, p\rangle = 0$, the line is tangent if $\langle p, q\rangle = 0$ as well, and we get a flex if $\langle q, q\rangle = 0$.

This line of reasoning only works if p is a smooth point on the curve, but we won't be talking about the special cases with singular points much in this class.

3 February 7, 2020

We'll talk today about the first subtle fact of this subject, the **dual curve**. Let $\mathbb{P} = \mathbb{P}^2$ be the projective plane, and consider a line L in $\mathbb{P}_x$ (that means points are labeled as (x_0, x_1, x_2)) with the equation $s_0x_0 + s_1x_1 + s_2x_2 = 0$.

We can think of points (x_0, x_1, x_2) on this line, but we can also think of this equation as a point (s_0, s_1, s_2) – both of these can be thought of as points in the projective plane, because the coefficients are only determined up to scaling. So that means that a line L in $\mathbb{P}$ gives a point (s_0, s_1, s_2), which we can call L^* in the **dual plane** $\mathbb{P}^* = \mathbb{P}^2_s$, and we can also take a point $p \in \mathbb{P}$ and turn it into a line p^* in $\mathbb{P}^*$ (by taking (a, b, c) and turning it into $ax_0 + bx_1 + cx_2 = 0$). One way we can then write down this interchange between points and lines is that

$$p \in L \iff L^* \in p^*$$

(the roles of points and lines switch).

So now let C be a plane curve of the form $f = 0$ (which is irreducible and homogeneous). Recall that a point $p \in C$ is **singular** if the partial derivatives $f_0(p) = f_1(p) = f_2(p) = 0$. All other points are **smooth** – it's easy to show that the number of singular points for an irreducible f is finite. Thus, let U be the set of smooth points; we can **map U to the dual space**

$$t : U \to \mathbb{P}^*.$$

Explicitly, how is this map defined? If we take a point $p \in U$, let L be the tangent line to C at p (that's why we leave out the singular points, because the tangent line isn't defined). Then we can let $t(p) = L^*$.

Theorem 7

Let $U^* = t(U)$. Then the closure of U^* is a curve C^* in the dual space $\mathbb{P}^*$, and $C^{**} = C$.

Basically, there's some special cases to think about for our curve C: maybe we have a **bitangent**, which means that a single tangent line is shared between two points, so the two points will have the same image in the dual space. (So the curve will cross itself, and this gives us a node in the dual space.) Also, maybe we have a **flex point** where the derivative goes to 0 for a moment – this turns out to give a cusp in the dual space. That's why we don't take

C to be smooth – C^* might still end up being singular even when C is smooth. Plus, we can use the degree of C to count the number of nodes and cusps, too, which can be interesting.

Example 8

A smooth cubic has 9 flex points (no bitangents because that would mean slightly perturbing the bitangent would make it intersect a line four times), and a generic quartic has 28 bitangents. (This means that the curve C^* coming from a generic quartic will have 28 nodes.)

How do we find the degree of the dual curve? If C is cubic and smooth, then C^* has degree 6, and it has 9 cusps (corresponding to the 9 flex points on the cubic). A famous picture for the quartic: intersect two ellipses, and take the product of the two equations to get an equation of the form

$$(ax^2 + by^2 - c)(a'x^2 + b'y^2 - c') = f(x, y).$$

And now if we add or subtract a small number, we can think about what happens to the zero locus. For example, the locus of zeros where $f(x, y) + \varepsilon = 0$ will look like "four beans" (in the four areas inside one ellipse but not the other), and that actually gives us 28 bitangents – four between each pair of beans, and one within each bean.

We're going to use something called the **transcendence degree**, which we should read on our own (it's algebra, not algebraic geometry). We'll review the main ideas here.

Definition 9

Let K/F be a field extension. $\alpha_1, \cdots, \alpha_k \in K$ are **algebraically dependent over F** if there exist coefficients $f(x_1, \cdots, x_k) \in F[x_1, \cdots, x_k]$ such that

$$f(\alpha_1, \cdots, \alpha_k) = 0.$$

(Otherwise, they are called **algebraically independent**.) The **transcendence degree** of K/F is the order of a maximal set of algebraically independent elements of K.

(It can be checked that the transcendence degree is independent of the choice of set that we use.)

Example 10

The field of rational polynomials $K = F(x_1, \cdots, x_k)$ has transcendence degree k, so any set of $k + 1$ elements is dependent.

With this, we can prove Theorem 7:

Proof. We're working in $K = \mathbb{C}(x_0, x_1, x_2)$ here, and we're looking for a polynomial $\phi(s)$ which is zero on $U^* = t(U)$. (Let's not worry about whether it's homogeneous or irreducible first.)

The transcendence degree of K over $\mathbb{C}$ is 3, so any four polynomials will be algebraically dependent: let's take $f(x)$ (which is the polynomial that defines C), as well as $f_0(x)$, $f_1(x)$, and $f_2(x)$ (the partial derivatives). Because these four polynomials are dependent, there exists an $\Psi(s_0, s_1, s_2, t)$ such that

$$\Psi(f_0(x), f_1(x), f_2(x), f(x)) = 0$$

is identically zero. We can assume Ψ does not have any factors of t, and now if we define $\phi(s_0, s_1, s_2) = \Psi(s_0, s_1, s_2, 0)$, then ϕ is now not the zero polynomial. On the other hand, if we let $\overline{x} = (\overline{x}_0, \overline{x}_1, \overline{x}_2)$ be a point of C (so $f(\overline{x}) = 0$

because we're on the zero locus), then

$$\phi(f_0(\overline{x}), f_1(\overline{x}), f_2(\overline{x})) = \Psi(f_0(\overline{x}), f_1(\overline{x}), f_2(\overline{x}), f(\overline{x})) = 0$$

(because Ψ is the zero polynomial here). So for all points $\overline{x}$ on U (that is, all smooth points of C of our curve), we have $\phi(f_0(\overline{x}), f_1(\overline{x}), f_2(\overline{x})) = 0$.

Now, what is the tangent line at this point $\overline{x}$? Remember the Taylor series expansion

$$f(p + tq) = f(p) + (\nabla_p q)t + O(t^2),$$

where $\nabla_p = (f_0(\overline{x}), f_1(\overline{x}), f_2(\overline{x}))$. So if $(s_0, s_1, s_2) = t(\overline{x})$ (where t is the map from U to $\mathbb{P}^*$), then $(s_0, s_1, s_2) = (f_0(\overline{x}), f_1(\overline{x}), f_2(\overline{x}))$. That means that $\phi(s) = 0$ is zero on U^*, and now we have a polynomial that vanishes on U^*. Since $\lambda\overline{x}$ is the same point as $\overline{x}$, we can make ϕ homogeneous (by just looking at one of its homogeneous parts). Now letting $g(x) = \phi(f_0, f_1, f_2)$, we know that g vanishes on U and is homogeneous: factor g, which corresponds to a factorization of ϕ. Now f will still divide one of the factors of g, and the corresponding irreducible factor ϕ will vanish on U (and therefore C): thus we've indeed shown that the closure of U^* is a curve in the dual space.

Let's spend some more time on why $C^{**} = C$. Say we have a point p_0 with tangent line L_0 on the curve C – this means that L_0^* will lie on p_0^* on the curve C^*. It seems pretty plausible that p_0^* will be the tangent line, and we'll try to show that here.

To do that, consider another point p_1 on C with tangent line L_1, and let L_0 and L_1 intersect at q. Since q lies on L_0 and L_1, the line q^* goes through points L_0^* and L_1^*. But now have p_1 approach p_0 – L_1^* will approach L_0^*, so the line q^* will approach the tangent line. And now we just need to show that q^* is equal to p_0^*.

We'll work with affine coordinates for this (specifically, work in $\mathbb{P}^2$ where $z = 1$) – pick them such that p_0 is the origin $(0, 0, 1)$ and the tangent line at p_0 is the line $y = 0$. Then C is some curve $f(x, y) = 0$ – we can solve for some function $y = y(x)$ such that $f(x, y(x)) = 0$. Letting $p_1 = (x_1, y_1)$, the equation of L_1 (the tangent to p_1) is of the form

$$y - y_1 = y_1'(x - x_1),$$

where $y_1 = y(x_1)$ and $y_1' = y'(x_1)$. Then $q = (x_q, 0, 1)$ lies on this tangent line, so we can solve to find

$$x_q = x_1 - \frac{y_1}{y_1'}.$$

As (x_1, y_1) goes to (x_0, y_0), x_q goes to 0, because the derivative y' has a zero of order one less than that of y. So the limit as p_1 goes to p_0 of q is indeed p_0, which is what we want – this tells us that we do have $C^{**} = C$. □

One quick application of this is that C being smooth means C^* has no bitangents and no flex points (otherwise C^*'s dual would have a node or cusp). Soon, we'll talk about the Plücker formulas, which tell us a bit more.

4 February 10, 2020

The first quiz has been moved from Friday to Wednesday due to popular vote.

Today, we're going to talk about **resultants**. Say we have two polynomials

$$f(x) = a_0x^3 + a_1x^2 + a_2x + a_3, \quad g(x) = b_0x^2 + b_1x + b_2$$

(they could also be homogeneous if we just put ys to bring the degrees up). Then the **resultant** Res(f, g) is a polynomial in the coefficients $\{a_i, b_j\}$, with the important property that $\text{Res}(f, g) = 0$ **if and only if f and g have a**

common root.

At first glance, it might be confusing why we care about such a polynomial – it turns out this actually computes a kind of "projection." Say that $f = f(t, x)$ and $g = g(t, x)$ are polynomials, so the coefficients a_i and b_j above are now polynomials in t. Then the resultant $\text{Res}_x(f, g)$ with respect to x is a polynomial in t as well – it has zeros at the "images" of the intersection points when we project down onto the t-axis.

We can do this in higher dimensions as well: say f and g are polynomials of x, y, z. Then $\text{Res}_z(f, g)$ is a polynomial in x and y, and it's zero at the points where we take the intersections of the zero loci for f and g and project that down onto the xy-plane.

So what polynomial in the coefficients $\{a_i, b_j\}$ are we trying to take here? Say f, g have a common root $x = \overline{x}$, and define $h = x - \overline{x}$. Then h divides f and g, so we can write $f = hp$ and $g = hq$ for polynomials p, q, so

$$\frac{fg}{h} = pg = fq$$

are two different ways of writing this polynomial. If f and g have, for example, degree 3 and 2, then p and q have degrees 2 and 1, respectively. We can think of $pg = fq$ as a relation among two polynomials – specifically, we can consider what monomials can appear in p. pg is a combination of x^2g, xg, and $1g$ if p is a quadratic, and fq is a combination of xf and $1f$. This gives us 5 **dependent polynomials**, and we'll put these in a matrix. Because the relation $pg = fq$ has degree 4, we'll write things in a 5×5 table.

	x^4	x^3	x^2	x	1
xf	a_0	a_1	a_2	a_3	0
$1f$	0	a_0	a_1	a_2	a_3
x^2g	b_0	b_1	b_2	0	0
xg	0	b_0	b_1	b_2	0
$1g$	0	0	b_0	b_1	b_2

We now have a 5 by 5 matrix R, and if the common root exists, then the **determinant of this matrix** should be 0 because the polynomials are dependent on each other.

Definition 11

The **resultant** of two polynomials f and g is the determinant of the matrix R that comes out of this process above.

Proposition 12

Suppose that f and g are monic, and suppose f has roots $\alpha_1, \alpha_2, \alpha_3$ and g has roots β_1, β_2. Then the resultant is the product

$$\text{Res}(f, g) = \prod_{i,j} (\alpha_i - \beta_j) = \prod_i g(\alpha_i)$$

(because $f = (x - \alpha_1)(x - \alpha_2)(x - \alpha_3)$ and $g = (x - \beta_1)(x - \beta_2)$).

This is not immediately obvious, but we'll show the proof. The idea is that the coefficients a_i and b_j are elementary symmetric polynomials in the roots α_i, β_j, so there must be a way to write this resultant in terms of a_i and b_j.

Proof. Let α_i, β_j be our roots, and let $f = \prod(x - \alpha_i)$ and $g = \prod(x - \beta_j)$. The resultant $\text{Res}(f, g)$ is a polynomial in the roots, and we divide this polynomial by $(\alpha_i - \beta_j)$, which is a monic polynomial in the α_i. Then

$$\text{Res}(f, g) = (\alpha_i - \beta_j)q + r,$$

where q, r are polynomials in α_i, β_j and r has degree 0 in α_i. Now if we set $\alpha_i = \beta_j$, $\text{Res}(f, g) = 0$ (because we have a common root), so $r = 0$ as well. But r can't depend on α_i, so $r = 0$, and therefore $\alpha_i - \beta_j$ divides the resultant for all i, j.

To finish, we just need to check that the degrees are the same – we claim that the resultant as a function of α_i and β_j is homogeneous with degree $\deg(f) \cdot \deg(g)$. (This would show that we're within a constant factor, and we can just use $f(x) = x^m$ and $g(x) = x^n - 1$ to check that constant factor. Then computing the determinant is not too bad – it's going to be ± 1.) We'll come back to that next time. □

For the rest of today, we'll talk about the **discriminant** of a polynomial.

Definition 13

The **discriminant** of a polynomial f is the resultant of f with its derivative f'.

For example, if $f(x) = ax^2 + bx + c$, then $f'(x) = 2ax + b$, and we have

$$\text{discr}(f) = \det \begin{bmatrix} a & b & c \\ 2a & b & 0 \\ 0 & 2a & b \end{bmatrix} = ab^2 + 4a^2c - 2ab^2 = 4a^2c - ab^2 = -a(b^2 - 4ac).$$

Often, people get rid of the a by working with a monic polynomial, and there's the extra negative sign, but the whole point is that this looks a lot like the $b^2 - 4ac$ we're familiar with.

What does the discriminant tell us? If f has a double root, then f' will share a root with f, and the discriminant will be zero. But also, if f is a function of t and x, we can plot it in the tx-plane – then the discriminant of f with respect to x is zero for values of t where there is a "vertical" tangent or singular point (because there's a double root in x at that fixed value of t).

Proposition 14

Let f be monic with roots $\alpha_1, \cdots, \alpha_n$. Then

$$\text{discr}(f) = \prod_{i \neq j} (\alpha_i - \alpha_j) = \prod_i f'(\alpha_i).$$

(Both (i, j) and (j, i) appear in the product.)

The logic here is the same as before – if the discriminant vanishes with $\alpha_i = \alpha_j$, then the discriminant as a polynomial in the roots must include $(\alpha_i - \alpha_j)$ as a factor for all i, j. For example, in the case where f is a quadratic,

$$f = (x - \alpha_1)(x - \alpha_2) \implies \text{discr}(f) = -(\alpha_1 - \alpha_2)^2.$$

As an application, we'll compute the genus of a smooth curve in the plane – let C be a smooth curve in $\mathbb{P}^2$ of degree d using coordinates (x, y, z), and we choose a generic point $q \in \mathbb{P}^2$. We'll project from $\mathbb{P}^2$ down onto $\mathbb{P}^1$: changing coordinates so that $q = (0, 0, 1)$, we draw lines from $(0, 0, 1)$ to points on the curve and see where they intersect $\mathbb{P}^1(x, y)$. Then the discriminant with respect to the variable z is 0 at the points p where the lines connecting q and p are tangent to the curve.

We'll assume for now that the points of tangency are double roots – this is a tricky point, and we'll talk about it next time. But the degree of the discriminant is $d^2 - d$ in the roots (take any pair of roots, and they're included twice in the product), so the projection has $d^2 - d$ **branch points**. Then the Euler characteristic of the curve is

$$e(C) = d \cdot e(\mathbb{P}^1) - (d^2 - d) = 2d - (d^2 - d) = 3d - d^2,$$

because the projective line $\mathbb{P}^1$ is just a sphere. Because the Euler characteristic $e(C) = 2 - 2g$, this tells us that the genus of a smooth curve of degree d is

$$g = \frac{1}{2}(d-1)(d-2).$$

Let's make a quick table:

d	1	2	3	4	5	6
g	0	0	1	3	6	10

This is interesting, because the genus of a curve can't be 2 or 4 or 5 if it's smooth. Also, let's compute the degree of the dual curve – the interesting points are where the lines from q is a tangent line to C. If we go over to the dual space $\mathbb{P}^*$ and look at our curve C^*, we have a line q^* which intersects it – the degree of C^* is just the number of intersections with the generic line q^*.

Well, every intersection L^* between C^* and q^* is some tangent line L to the original curve C, so the degree of the dual curve is the same as the degree of the discriminant! Thus, the degree of C^*, the dual curve, is equal to $d^2 - d$. (And this verifies that the degree of the dual curve for a generic cubic is 6.)

5 February 12, 2020

The due date for problem set 1 was pushed back to Friday, because we'll cover a few more ideas today that will be interesting.

Say that we have an affine curve C corresponding to $f(x, y) = 0$, and we separate out the different degrees via

$$f = f_0 + f_1 + \cdots + f_d.$$

Let $p = (0, 0)$ – then $p \in C$ if $f_0 = 0$, and p is a **smooth point** of C if $f_0 = 0$ and $f_1 \neq 0$. Meanwhile, if $f_0 = f_1 = 0$ but $f_2 \neq 0$, then p is called a **double point** – in general, a point has **multiplicity** k if $f_0, \cdots, f_{k-1}$ are all zero, but $f_k(p) \neq 0$.

We'll look at specifically at double points: then we can write

$$f = (ax^2 + bxy + cy^2) + (dx^3 + \cdots).$$

We'll analyze by "blowing up" the xy-plane: let $t = \frac{y}{x}$, which means that we substitute in $y = tx$. Then we have a map π from the xt-plane to the xy-plane given by $(x, y) = (x, tx)$: the t-axis gets sent to the origin under π. So this is a blowup because the origin gets sent to a line, and the y-axis doesn't get hit except at the origin, but everything else goes bijectively. (And horizontal lines in the xt-plane go to lines through the origin.)

So if we substitute $y = tx$ into our polynomial f, we get a polynomial (we can get rid of the x^2 factor throughout)

$$g(x, t) = \frac{f(x, tx)}{x^2} = a + bt + ct^2 + dx + \cdots,$$

where everything else after the dx term is divisible by x. The constant term (if we look at g as a function of x) is $q(t) = a + bt + ct^2$. **If** $q(t)$ **has distinct roots**, and we normalize our function so that $c = 1$ (when it's 0, just change the coordinates slightly), we have

$$g(x, t) = (t - \alpha)(t - \beta) + dx + \cdots.$$

Then the partial derivative with respect to t is

$$g_t = 2t - \alpha - \beta + x(\cdots).$$

If we look at the points in the (x, t)-plane where g is zero, we have $(0, \alpha)$ and $(0, \beta)$, so the curve goes through those two points. But $g_t(0, \alpha) = \alpha - \beta \neq 0$, and we can solve $g(x, t) = 0$ to find an analytic function $t = u(x)$ for $v(0) = \alpha$ (for small x). Similarly, we can find an analytic function $t = v(x)$ with $v(0) = \beta$ – the point is that the projection π to the xy-plane gives us two lines through the origin. This means the curve intersects itself, which is called a **node**.

What about the case where $q(t)$ has a repeated root? Then g looks like

$$g(x, t) = (t - \alpha)^2 + x(d + \cdots).$$

Now $g_t(0, \alpha) = 0$, but the partial with respect to x, $g_x(0, \alpha)$ is nonzero **if** $d \neq 0$. So we can again solve $x = w(t)$ for small t: it'll have a vertical tangent at $\alpha = 0$ in the xt-plane. Projecting this back down to the xy-plane gives us a **cusp**. And the last case is where we have a double root, and we also have $d = 0$: then the locus $g = 0$ is singular, so the conclusion is that the blowup curve $g = 0$ is smooth above a point p if and only if the singularity is a node or a cusp.

Let's look a bit more at the geometry of the cusps – the **standard cusp** is the point $(0, 0)$ for the curve $y^2 - x^3 = 0$. All cusps are analytically equivalent to this one, so we'll look a bit at the geometry here. We parameterize via $x = t^2$ and $y = t^3$: the question we're asking is "what does the cusp look like if we slice across a cross-section?"

To figure this out, let $t = e^{i\theta}$. Then $x = e^{2i\theta}$ and $y = e^{3i\theta}$, so as θ ranges from 0 to 2π, x goes around the unit circle twice, while y goes around three times. The product of the two unit circles is a torus, so this gives us a picture on the surface of the torus traced out by the cusp as seen below:

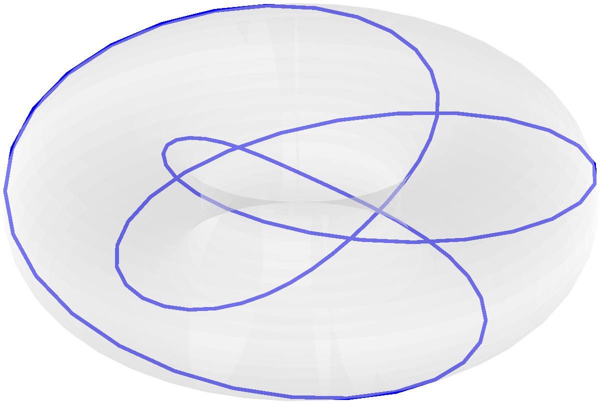

If we look at this more closely, it's actually a trefoil knot! So that's what a cusp looks like on the cross-section $|x| = |y| = 1$. (But in four real dimensions, it doesn't really make sense to call it a "knot" anymore.) We could also think about the unit sphere $x\overline{x} + y\overline{y} = 1$ and intersect it with a cusp, and do a stereographic progression from the 3-sphere into 3-space – this also gives a kind of knot.

We'll now move on: **Hensel's lemma** will be our next topic. Let $f(x), g(x)$ be polynomials; we can write out the polynomial $p(x) = f(x)g(x)$:

$$f(x)g(x) = (a_0x^3 + a_1x^2 + a_2x + a_3)(b_0x^2 + b_1x + b_2) = c_0x^5 + c_1x^4 + c_2x^3 + c_3x^2 + c_4x + c_5 = p(x).$$

The coefficients c_k can be computed in terms of the a_i and b_j:

$$\begin{array}{llllll} c_0 = & a_0 b_0 & & & & \\ c_1 = & a_0 b_1 & +a_1 b_0 & & & \\ c_2 = & a_0 b_2 & +a_1 b_1 & +a_2 b_0 & & \\ c_3 = & & a_1 b_2 & +a_2 b_1 & +a_3 b_0 & \\ c_4 = & & & a_2 b_2 & +a_3 b_1 & \\ c_5 = & & & & a_3 b_2. & \end{array}$$

These are the "product equations," and it's an interesting question to ask whether we can solve for the a_i and b_j. One useful idea is the **Jacobian criterion**: say that f and g are monic, so that $a_0 = b_0 = 1$. Then let the **Jacobian matrix** be

$$J = \frac{\partial(c_1, c_2, c_3, c_4, c_5)}{\partial(b_1, b_2, a_1, a_2, a_3)}.$$

We can calculate the entries of this matrix explicitly:

$$J = \begin{bmatrix} a_0 & 0 & b_0 & 0 & 0 \\ a_1 & a_0 & b_1 & b_0 & 0 \\ a_2 & a_1 & b_2 & b_1 & b_0 \\ a_3 & a_2 & 0 & b_2 & b_1 \\ 0 & a_3 & 0 & 0 & b_2 \end{bmatrix}.$$

This is the transpose of the resultant matrix R. Thus, the Jacobian matrix is singular (has determinant 0) if and only if f and g have a common root, which gives us the following result:

Corollary 15 (Hensel's lemma)

Suppose that the coefficients a_i, b_j are analytic functions of t, and let $\overline{a}_i = a_i(0), \overline{b}_j = b_j(0)$. Here $f = f(t, x)$ and $g = g(t, x)$ are polynomials in x with coefficients that are polynomials in t, and $p = fg$: let $\overline{f} = f(0, x)$ and define $\overline{g}$, $\overline{p}$ similarly. Then if $\overline{f}, \overline{g}$ have no common roots, and $\overline{f}$ is monic, then there exist f, g such that

$$p(t, x) = f(t, x)g(t, x)$$

and $\overline{p} = \overline{f}\,\overline{g}$ for small t.

Proof. Since $\overline{f}$ is monic, we have $a_0 = 0$, so the leading coefficient b_0 for the other factor is also nonzero. Now the Jacobian matrix is nonsingular, so the Implicit Function Theorem tells us that there is a unique solution for the remaining coefficients for small t, as desired. (If we don't know what the Implicit Function Theorem is, we should read this on our own. Here, we're treating different powers of x as different "functions.") □

Basically, if we have some curve $p(t, x) = 0$ and we're interested in how it looks near $t = 0$, suppose that $p(0, x)$ factors into two polynomials with distinct roots. Then Hensel's lemma says that we can factor $p(t, x)$ near 0.

The simplest example is if we have two polynomials

$$c_0(t)x^2 + c_1(t)x + c_2(t) \overset{?}{=} (x + a_1)(b_0 x + b_1),$$

where a_1, b_0, b_1 are functions of t. The product equations are

$$c_0 = b_0, c_1 = a_1 b_0 + b_1, c_2 = a_1 b_1.$$

We're given that $\overline{p} = p(0, x)$ factors as $(x + \overline{a}_1)(\overline{b}_0 x + \overline{b}_1)$, and those two linear factors don't have a root in common.

This means that $\overline{a}_1 \neq \frac{\overline{b}_1}{\overline{b}_0}$, and Hensel's lemma tells us that p factors via

$$p(t,x) = (x + a_1(t))(b_0(t)x + b_1(t)).$$

6 February 14, 2020

We'll talk about the Plücker formulas today: they count things like bitangents, nodes, and cusps. The formulas are completely unimportant, but the interesting thing is that there are formulas (and they don't depend on anything except for degrees).

Specifically, we'll work with a degree d smooth curve C which is **ordinary**, which means the following things:

- Flex points are ordinary (we have a triple intersection but not higher).
- Bitangents are ordinary – neither point is a flex point, and there are no tritangents.

This means that our map $t : C \to C^*$ from our curve to the dual curve is defined everywhere. Pick a point not on the curve, and choose coordinates so that $q = (0,0,1)$. For every point p on the curve, draw a line from q to p, and project this down to the xy-plane.

What does this look like in the dual plane? We have a special line q^* which will intersect our curve C^*, and it intersects at points L^* corresponding to **tangent lines** in our original plane. In other words,

$$\deg C^* = \#(C^* \cap q^*),$$

and this is the number of points $\tilde{p}$ which are images of tangent lines from q.

We can calculate this using the discriminant: at each tangent line from q, the discriminant should vanish, because we have a repeated root. So the degree of C^* is the number of zeros of the discriminant f with respect to one of the variables z. All of these zeros are **simple zeros**, and the discriminant of f with respect to z is the resultant

$$\text{Res}\left(f, \frac{\partial f}{\partial z}\right).$$

f has degree d, and $\frac{\partial f}{\partial z}$ has degree $d-1$, so the resultant has degree $\boxed{\deg C^* = d(d-1)}$ – this is our first result.

We have some other numbers that we care about – the number of flex points f of C, as well as the number of bitangents b. First, we can count f:

Lemma 16

A smooth point p of our curve C is a flex point if and only if the determinant of the Hessian matrix H_p (at the point p) is 0, where $H_p = \left(\frac{\partial^2 f}{\partial x_i x_j}\right)_{1 \le i,j \le 3}$.

Proof. Remember that we defined the bilinear form

$$\langle u, v \rangle = u^t H_p v$$

where we're looking at the points as column vectors. Then we know that p is a flex if and only if the first three terms of the Taylor expansion

$$f(t + pq) = \frac{1}{d(d-1)}\langle p, p\rangle + \frac{1}{d-1}\langle p, q\rangle t + \frac{1}{2}\langle q, q\rangle t^2 + O(t^3)$$

vanish. If f is a flex point, then the form must be degenerate, because it's operating on a three-dimensional vector space where p and q are independent vectors. On the other hand, if the determinant is zero, there is a null vector, so there exists q which satisfies this equation above. □

With this, how can we count the flex points? Let H be the **Hessian curve** $\{p : \det H_p\} = 0$: we wish to compute the degree of H. H is a 3 by 3 matrix, and all entries have degree $d-2$. This means that the determinant $\det H_p$ has degree $3(d-2)$.

Here, we're going to use a useful result:

Theorem 17 (Bézout)
Let X and Y be curves of degree m and n in projective space. Then the number of intersections is mn if intersections are transversals (not double multiplicity).

(We'll prove this using cohomology later on in the class – it's not a very deep theorem, but it's interesting.) Also, we'll use a result which is ugly to prove (mostly just computation):

Lemma 18
If p is an ordinary flex of our curve C, then C intersects the Hessian divisor curve H transversely at p.

So now, if the degree of H is $3(d-2)$, and the degree of C is d, the number of flexes is just the number of intersections:

$$f = 3d(d-2) = 3d^2 - 6d$$

(as long as $d \geq 2$). It's a fact that flexes map to cusps in the dual curve, and bitangents map to nodes, but this isn't obvious. (We'll assume it for now.) Define δ^* to be the number of nodes of C^* and κ^* to be the number of cusps: then $\boxed{\delta^* = b, \kappa^* = f}$.

Counting the number of bitangents is a bit harder, and there isn't a very easy way to directly compute them on our curve C. So instead, let's count the number of nodes on C^*! Remember that the bidual gives us back the original curve: $(C^*)^* = C$, as long as we take closures. Choose a generic point Q in the dual plane, and project the curve from Q onto the xy-plane (still in the dual plane). Then there are three types of interesting lines: (1) tangent lines to our dual curve, (2) lines through a node, and (3) lines through a cusp.

Remember that we have some equation $\phi(u, v, w) = 0$ for our dual curve, and the discriminant with respect to one of the variables v has simple zeros for case (1), double zeros for case (2), and triple zeros for case (3). And this gives us a formula: we know the degree of C^* is $d^* = d(d-1)$, so the degree of the discriminant is

$$\deg \operatorname{discr}_v(\phi) = d^*(d^*-1).$$

How many tangents do we have (case (1))? Remember the tangent lines correspond to intersections in the dual with our generic line q^*. So because we're working with the dual curve here, we should have $d = \deg C$ total tangent lines. So our equation is now

$$d^*(d^*-1) = 1 \cdot d + 2 \cdot \delta^* + 3 \cdot \kappa^*.$$

But we know d^* and κ^*, so this allows us to compute $\delta^* = b$: substituting in $\delta^* = b$ and $\kappa^* = f = 3d^2 - 6d$, we have that

$$(d^2-d)(d^2-d-1) = d + 2b + 3(3d^2-6d).$$

Expanding this out tells us that

$$d^4 - 2d^3 - 9d^2 + 18d = 2b.$$

This is zero for $d = 3$ – cubics can't have bitangents, so the first interesting case is $d = 4$:

$$256 - 128 - 144 + 72 = 2b \implies b = 28,$$

which is the result we've been talking about for a while.

So how do we show that there are double zeros for nodes and triple zeros for cusps? The main idea is to use Hensel's lemma. Remember that if we have a curve $f(x, y) = 0$ in the plane and we project a node down onto the x-plane, we get a double zero. Say that the node is at $(0, 0)$: then we know that

$$\overline{f}(y) = f(0, y) = y^2\overline{h}(y)$$

for some other polynomial $\overline{h}$. (The other zeros of h come from the other intersections of $x = 0$ with our curve $f = 0$.) Then Hensel's lemma tells us that we can also factor the multivariable polynomial

$$f(x, y) = g(x, y)h(x, y)$$

for small x, where $g(0, y) = y^2$ and $h(0, y) = \overline{h}$. What do we know about the discriminant of f? Treating x as a parameter,

$$\text{discr}(g(y)h(y)) = \prod_{i,j}(\alpha_i - \alpha_j)\prod_{k,\ell}(\beta_k - \beta_\ell)\prod_{i,k}(\alpha_i - \beta_k).$$

where α_i are the roots of g and β_k are the roots of h. But the first two terms are the discriminants of g and h, and the last term is the resultant. So

$$\text{discr}(f) = \text{Res}(g, h)\text{discr}(g)\text{discr}(h).$$

So if we avoid accidents, the discriminant of the polynomial $\overline{h}$ and the resultant of g and h will be nonzero at $x = 0$. So the order of vanishing of g is the same as the order of vanishing of f, and then we compute things for a quadratic using the usual formulas.

Starting next class, we'll move on to the second chapter – basic algebraic geometry for affine varieties.

7 February 18, 2020

We'll start by talking about the **Zariski topology** today on the affine space $\mathbb{A}^n = \{x = (x_1, \cdots, x_n) : x_i \in \mathbb{C}\}$. Denote $\mathbb{C}[x] = \mathbb{C}[x_1, \cdots, x_n]$, and let $f = (f_1, \cdots, f_k)$ be a set of polynomials $f_i \in \mathbb{C}[x]$. Then we'll define

$$V(f) = \{x \in \mathbb{A}^n | f(x) = 0, \text{ i.e. } f_i(x) = 0 \ \forall i\}.$$

These are known as the **Zariski closed sets**. There are a few things to understand about them, none of which are particularly difficult to prove: for example, if I is the ideal generated by f, often denoted (f) (remember this is vector notation), then I consists of

$$I = \{g_1f_1 + \cdots + g_kf_k, \quad g_i \in \mathbb{C}[x]\}.$$

Then it's clear that if the f_is are zero at a point, then everything in the ideal is also zero at that point, so $V(f) = V(I)$. (The Basis Theorem tells us that every ideal in $\mathbb{C}[x]$ will look like $(f_1, \cdots, f_k)$.) Also, we have a nesting property: if $I \supset J$, then $V(I) \subset V(J)$. (The inclusions are reversed because having more functions is more restrictive, so there are less points in $V(I)$.)

Proposition 19

The Zariski closed sets are closed in the Zariski topology.

Proof. We need to check that $\varnothing$ and $\mathbb{A}^n$ are Zariski closed, and if C_j are closed, the (possibly infinite) intersection $\bigcap C_j$ is Zariski closed. We also need to show that $C \cup D$ is closed (which shows this for finite unions as well). All of these are pretty obvious except the last one.

$\varnothing$ is the zero locus $V(1)$ (it's where $1 = 0$) and $\mathbb{A}^n$ is the zero locus $V(0)$ (it's where $0 = 0$). To show that intersections are closed, say that $C_j = V(I_j)$: then

$$\bigcap_j C_j = V\left(\sum I_j\right),$$

because the intersection of the zero loci are where all the polynomials in each of the I_js are zero. Since $\sum I_j$ is also an ideal, $\bigcap C_j$ is also Zariski closed, as desired.

The last one – showing $C \cup D$ is closed – is slightly nontrivial: we'll need to notice something about two ideals, which is that

$$I \cap J \supset IJ \supset (I \cap J)^2.$$

(Here, IJ is the product ideal, which is the **sums of** products of elements of I and J.) From these inclusions, we can see that if $C = V(I)$ and $D = V(J)$, then

$$V(I \cap J) \subset V(IJ) \subset V((I \cap J)^2),$$

but the zero locus of an ideal squared is the same as the zero locus of the ideal. Therefore, we actually have

$$V(I \cap J) = V(IJ).$$

With this, we want to show that $C \cup D = V(IJ)$. We know that $I \supset IJ$, so $V(IJ) \supset C$, and similarly $V(IJ) \supset D$, so this means $V(IJ) \supset C \cup D$.

To show the other direction, we want to show that if $x \in V(IJ)$, then $x \in C$ or $x \in D$. We know that for all $f \in I$ and $g \in J$, if $x \in V(IJ)$, then $f(x)g(x) = 0$. If $f(x) = 0$ for all $f \in I$, then $x \in C$. Otherwise, there exists an f such that $f(x) \neq 0$. Using this f and looking over all $g \in J$, we see that we must have $g(x) = 0$ for all g, meaning $x \in D$. This shows that $V(IJ) \subset C \cup D$, and we're done. □

Remark 20. *It's important to keep in mind that closed sets in the Zariski topology look very different from what we usually call "closed sets."*

One nice property of $\mathbb{C}[x]$ is that it is **Noetherian**: any increasing chain of ideals $I_1 < I_2 < \cdots$ must be finite. What does this tell us about the Zariski closed sets? Let C_j be Zariski closed sets, and let $I_j = \{f | f = 0 \text{ on } C_j\}$. This is an ideal, and $C_1 > C_2$ if and only if $I_1 < I_2$. This gives us an analogous condition to the one on ideals:

Proposition 21

The Zariski topology has the **descending chain condition** on closed sets: if $C_1 > C_2 > \cdots$ is a decreasing chain, it must be finite.

Zariski open sets are the complement of Zariski closed sets, so we get an ascending chain condition for open sets as well.

Definition 22

A topological space is **Noetherian** if it has a descending chain condition on closed sets.

If X is a topological space and S is a subset of X, then S becomes a topological space with the **induced topology**: $T \subset S$ is closed if $T = S \cap C$ and C is a closed set in X. (And restricting to subsets of X will still give us the descending chain condition.)

Definition 23

A topological space X is **irreducible** if it is not the union of two proper closed subsets.

In other words, if C and D are closed in an irreducible space X, and $X = C \cup D$, then $C = X$ or $D = X$.

Proposition 24

Let X be an irreducible topological space. If $U \subset X$ is a (nonempty) open set, then the **closure** of U (smallest closed set containing U) is X.

Proof. Say that $\overline{U}$ is the closure of U. Since X is irreducible, if we let $C = X - U$ be the complement of U,

$$X = U \cup C \implies X = \overline{U} \cup C.$$

And now C is not X, because that means U is nonempty, so $\overline{U}$ must be X. □

Definition 25

A topological space X is **connected** if it is not the union of two proper disjoint closed sets.

It's hard to satisfy irreducibility, because we don't require our two closed sets C and D to be disjoint! For example, two intersecting lines form a connected space but not an irreducible space.

Theorem 26

Let X be a Noetherian space. Then X is a finite union of irreducible spaces.

Proof. We prove the contrapositive. Suppose that $X = C_0$ is not a finite union of irreducible spaces: then we can write $C_0 = C_1 \cup D_1$, where C_1, D_1 are proper closed subsets of C_0. Then one of C_1 and D_1 is not a union of finitely many irreducible spaces: say that C_1 is not. Then we know that $C_0 > C_1$, and now we repeat this argument with C_1 instead of C_0. But this gives us an infinite descending chain of closed sets, so X is not Noetherian. □

Definition 27

A Zariski closed subset X of $\mathbb{A}^n$ is a **variety** if it is irreducible.

We'll ask a question (which we'll answer next time): if I and J are ideals, when is $V(I) = V(J)$? It's clear that I and J don't need to be equal: I can be generated by f and J can be generated by f^2. But it turns out that this is really the only reason that two ideals would have the same variety:

Definition 28

The **radical** of an ideal I, denoted $\operatorname{rad} I$, is the ideal

$$\operatorname{rad} I = \{f : f^n \in I \text{ for some } n > 0\}.$$

$\operatorname{rad}(I)$ contains I, and the zero loci are the same: $V(\operatorname{rad} I) = V(I)$ (because f is zero if and only if $f^n = 0$). This radical is the key to answering our question, but it'll take some work to prove:

Theorem 29

$V(I) \supset V(J)$ if and only if $\operatorname{rad} I \subset \operatorname{rad} J$.

The backwards direction is clear – if $\operatorname{rad} I$ is contained in $\operatorname{rad} J$, then $V(\operatorname{rad} I)$ contains $V(\operatorname{rad} J)$, so $V(I)$ contains $V(J)$. But the other direction is something we'll show tomorrow.

Another question we may want to ask: when is a Zariski closed set a variety?

Theorem 30

$X = V(I)$ is a variety (in other words, it is irreducible) if and only if $\operatorname{rad}(I)$ is a **prime ideal**.

Recall that an ideal P is **prime** if

$$ab \in P \implies a \in P \text{ or } b \in P.$$

Alternatively, if A, B are ideals,

$$AB \subset P \implies A \subset P \text{ or } B \subset P.$$

Let's do an example:

Example 31

Consider a map ϕ from $\mathbb{A}^2_{x,y}$ to $\mathbb{A}^4_{t,u,v,w}$ of the form

$$\phi(x, y) = (x^3, y^3, x^2y, xy^2).$$

There are some relations here: notice that $tu = vw$, $v^3 = t^2u$, and $w^3 = tu^2$. These three relations generate some ideal I. Let's show that the radical $\operatorname{rad} I$ is a prime ideal: we want to show that the locus of zeros $X = V(I)$ here is irreducible.

To do this, we'll show that ϕ maps $\mathbb{A}^2$ surjectively to X. Given t, u, v, w which satisfy these relations, we just need to find x and y: we can just set $x = t^{1/3}, y = u^{1/3}$. There are three choices for each of x and y, and we just need to check that $\phi(x, y) = (t, u, v, w)$ – as long as our relations are satisfied, there is a choice of x and y that work here.

With this, it is easy to show that X is irreducible. Suppose otherwise: then its inverse image is the union of two closed sets. But $\mathbb{A}^2$ is irreducible, which is a contradiction! So we have indeed shown that $\operatorname{rad} I$ is a prime ideal. And in this particular case, $\operatorname{rad} I = I$, but it requires more work to show this.

8 February 19, 2020

We'll start with a bit of review: in the Zariski topology on the affine space $\mathbb{A}^n$, the closed sets are of the form $V(I)$ for an ideal I in $\mathbb{C}[x_1, \cdots, x_n]$: they are the points p which vanish for all polynomials $f \in I$. (By the Basis Theorem, all such ideals I are finitely generated.)

We know that $\mathbb{A}^n$ is a **Noetherian topological space** – basically, it has the descending chain condition on closed sets, which follows from the ascending chain condition on ideals. Thus, every closed set can be decomposed into a finite union of irreducible closed sets. (Recall that X is **irreducible** if it cannot be broken up into proper subsets $C \cup D$.) Then an **affine variety** is an irreducible closed set under the Zariski topology.

At the end of last class, we defined the **radical**

$$\text{rad}\, I = \{g : g^n \in I \text{ for some } n \geq 1\}.$$

We know that $I \subset \text{rad}\, I$, and also

$$V(I) = V(\text{rad}\, I)$$

because a polynomial f has the same zeros as f^n for any positive integer n. We'll be proving some things today that yield the following results:

Proposition 32

We have

- $V(I) \supset V(J)$ if and only if $\text{rad}(I) \subset \text{rad}(J)$.
- Let P be a radical ideal (meaning $\text{rad}\, P = P$). Then $V(P)$ is a variety if and only if P is a prime ideal.

Here's the first main result:

Theorem 33 (Hilbert Nullstellensatz)

Let $\mathbb{C}[x] = \mathbb{C}[x_1, \cdots, x_n]$. There is a bijective correspondence between the following three sets:

- Points p of affine space $\mathbb{A}^n$,
- Homomorphisms π_p from $\mathbb{C}[x]$ to $\mathbb{C}$,
- Maximal ideals $\mathfrak{m}_p$ of $\mathbb{C}[x]$.

Basically, π_p evaluates a polynomial $f \in \mathbb{C}[x]$ at p, and $\mathfrak{m}_p$ is the kernel of π_p.

This comes from 18.702, so we won't go into it in much detail. For example, it's clear that every homomorphism is surjective, so the kernel is a maximal ideal. If we're looking at a point $p = (a_1, \cdots, a_n)$, then the maximal ideal $\mathfrak{m}_p$ is generated as

$$\mathfrak{m}_p = (x_1 - a_1, \cdots, x_n - a_n).$$

Definition 34

An **algebra** A is a ring that contains the complex numbers $\mathbb{C}$. A **finite-type** algebra A can be generated as an algebra by some finite set of elements $\{\alpha_1, \cdots, \alpha_n\}$, where $\alpha_i \in A$, so that all elements can be written as polynomials in α_i with coefficients in $\mathbb{C}$.

Another way to define a finite-type algebra is to consider the map

$$\tau : \mathbb{C}[x_1, \cdots, x_n] \to A$$

which sends x_i to α_i. Then this map should be **surjective** if we can write every element of our algebra as a polynomial in the α_is, and thus we can write

$$A \cong \mathbb{C}[x_1, \cdots, x_n]/I$$

for some ideal $I = \ker \tau$ by the First Isomorphism Theorem.

Fact 35

A finite module over a ring means that every element in the module is a linear combination of our generating set, but a finite-type algebra is different – it lets us write general polynomials! So we should be careful not to confuse the two.

Theorem 36 (Hilbert Nullstellensatz, version 2)

Let A be a finite-type algebra. Then there is a bijective correspondence between homomorphisms $\overline{\pi} : A \to \mathbb{C}$ and maximal ideals $\overline{\mathfrak{m}}_p$ of A.

Proof. This is basically the first version but applying the Correspondence Theorem. We can present

$$A = \mathbb{C}[x_1, \cdots, x_n]/I,$$

and then our homomorphisms $\overline{\pi} : A \to \mathbb{C}$ correspond bijectively to homomorphisms $\pi : \mathbb{C}[x] \to \mathbb{C}$ where $\ker \pi$ contains the ideal I. (And the rest of the argument follows analogously.) □

So if $A = \mathbb{C}[x]/I$, the two sets in the above theorem (maximal ideals of A and homomorphisms to $\mathbb{C}$) also correspond bijectively to the **locus of zeros of** I in $\mathbb{A}^n$. We'll actually use this more: A has to be a finite-type algebra, but we don't have to express it explicitly in the form $\mathbb{C}[x_1, \cdots, x_n]/I$ for the result to hold. The point is that we can skip the presentation altogether:

Definition 37

Let A be a finite-type **domain** (no zero divisors). The **spectrum** of A, denoted $\operatorname{Spec} A$, is a set of points where we put p into $\operatorname{Spec} A$ for every maximal ideal $\mathfrak{m} = \mathfrak{m}_p$ of A.

By definition, then, the set of points of $\operatorname{Spec} A$ correspond to the maximal ideals $\overline{\mathfrak{m}}_p$, which correspond to the homomorphisms $\overline{\pi}_p$. So if we have a presentation $A = \mathbb{C}[x_1, \cdots, x_n]/I$, then $\operatorname{Spec} A$ is just the zeros of the ideal $V(I)$.

Theorem 38 (Strong Nullstellensatz)

Let $f_1, \cdots, f_k, g \in \mathbb{C}[x_1, \cdots, x_n]$, and let $V = V(f)$ be the set of zeros of (all of) $f_1, \cdots, f_k$ in $\mathbb{A}^n$. Then if g is identically zero on V, then some power of g is in the ideal $I = (f_1, \cdots, f_k)$:

$$g^n = h_1 f_1 + h_2 f_2 + \cdots + h_k f_k,$$

where $h_i \in \mathbb{C}[x]$.

Proof by Rainich. Add another variable y to the polynomial ring to get $\mathbb{C}[x_1, \cdots, x_n, y]$, and let W be the locus of zeros in $\mathbb{A}^{n+1}_{x,y}$ where $f_1 = f_2 = \cdots = f_k = 0$ and $gy - 1 = 0$. (In other words, $g(x_1, \cdots, x_n)y = 1$.) Consider some $(x, y) \in W$ in the zero locus: then

$$f_1(x) = \cdots = f_k(x) = 0$$

implies that $g(x) = 0$ as well (because g is identically zero on V). But then we can't solve $gy = 1$, so **the locus** W **is empty**. Now, what can we say about a set of polynomials (specifically $(f_1, \cdots, f_k, gy - 1)$ here) whose zero locus is empty? We'll use a quick lemma here:

Lemma 39

Any ideal $I < R$ of a ring R is contained in a maximal ideal.

Using that fact, note that maximal ideals correspond to points in $\mathbb{A}^n$, so if there are no points in the zero locus, the ideal I generated by $f_1, \cdots, f_k, gy - 1$ must be the **whole ring** R (it's not contained in a maximal ideal). Thus, there exist $p_1(x, y), \cdots, p_k(x, y), q(x, y) \in \mathbb{C}[x, y]$ which satisfy

$$p_1 f_1 + \cdots + p_k f_k + q(gy - 1) = 1.$$

Now let's go to the ring $R = \mathbb{C}[x][y]/(gy-1)$: this means that the residue of y is g^{-1}. (In other words, we've adjoined an inverse of $g(x)$ to $\mathbb{C}[x]$.) We can suppose $g \neq 0$ (or the statement is trivial), and thus $\mathbb{C}[x] \subset R$.

But we also have our relation equation above, so in our ring R, $gy - 1 = 0$. This means that in R, we have

$$p_1 f_1 + \cdots + p_k f_k = 1.$$

The p_is are polynomials in x and y, but we can get rid of $y = g^{-1}$ by multiplying by a sufficiently high power of g. So if we multiply by g^N for some N and cancel the y's, we have a relation

$$h_1(x) f_1(x) + \cdots + h_k(x) f_k(x) = g(x)^N,$$

which is exactly what we want – this means g^N is in the ideal generated by f, as desired. □

The first two results that we stated today, Proposition 32, can now be deduced from the Strong Nullstellensatz pretty easily – this is an exercise!

The idea here is that we're **localizing** our ring, which means to adjoin an inverse. What's the meaning of this name? In topology, something is true locally if it's true in some open neighborhood. There are some open sets that we understand well, and they are important – this proof illustrates some of the power of this startegy here!

9 February 21, 2020

We'll start with a bit of review: recall that a **finite domain** A is isomorphic to some $\mathbb{C}[x]/P$, where P is a prime ideal. We define $X = \operatorname{Spec} A$ to contain the points p corresponding to maximal ideals $\mathfrak{m}_p$, which correspond to homomorphisms $\pi_p : A \to \mathbb{C}$. Also, it's good to remember that $\operatorname{Spec} A$ for a domain $A = \mathbb{C}[x]/I$ corresponds to the variety $V(I)$.

Elements $\alpha \in A$ determine functions on $X = \operatorname{Spec} A$ via the correspondence between points p in the zero locus and maximal ideals $\overline{\pi}_p$:

$$\alpha(p) = \overline{\pi}_p(\alpha).$$

These elements of A are called the **regular functions** on X.

Lemma 40

The function $\pi_p(\alpha)$ determines the element α. Equivalently (because π_p is a homomorphism), if a function f is zero, then $\alpha = 0$.

Proof. If $A = \mathbb{C}[x]/P$, then $X = V_{\mathbb{A}^n}(P)$ in affine space $\mathbb{A}^n$ is the zeroset of the ideal P. (Note that the correspondence theorem tells us that maximal ideals of A are the same as maximal ideals of $\mathbb{C}[x]$ that contain P.) So a function f

that is zero on X is in P, so it must be the zero element α in A. □

Definition 41

In the **Zariski topology** on an affine variety X, a closed set is the set of zeros of some ideal I of A.

For the sake of notation, denote $V_X(I)$ to be the zeroset of I: $p \in V_X(I)$ means that all elements $\alpha \in I$ are zero at p: this means that α is in the kernel of π_p, and we can correspond this to the maximal ideals via

$$V_X(I) = \{p : I \subset \mathfrak{m}_p\}.$$

One important operation which we briefly talked about last time is the idea of **localization**. Let A be a finite domain, and let $X = \operatorname{Spec} A$. Pick a nonzero element $s \in A$ – we can adjoin an inverse to A, which we denote as

$$A_s = A[s^{-1}] = A[z]/(zs - 1).$$

A_s is called the **localization** of A, and $X_s = \operatorname{Spec} A_s$ is similarly called the localization of X.

Proposition 42

X_s corresponds to points of X at which $s(p) \neq 0$, and it's an open subset of X because it's the complement of $s = 0$. Then the Zariski topology on X_s is the induced topology (a closed set of X_s is a closed set in X, intersected with X_s).

Proof. Let $p \in X$ be a point, and we correspond this to the homomorphism $\pi_p : A \to \mathbb{C}$. If $s(p) \neq 0$, then we can extend π_p to a homomorphism $A_s \to \mathbb{C}$ in a unique way: we know where the elements of A go, and the inverse s^{-1} goes to the inverse of $s(p)$, which determines the homomorphism. On the other hand, if $s(p) = 0$, we can't define where the inverse goes. So whenever we have a homomorphism π_p, we have a point $p \in \operatorname{Spec} A_s$.

For the topology, let D be a closed set in X_s: this means D is the set of zeros of some set of elements of A_s. But an element of A_s looks like as^{-n}, where $a \in A$, so we have some set

$$\{a_1 s^{-n}, \cdots, a_k s^{-n}\}$$

of polynomials. (We can use the same exponent for all of the a_i.) This has the same zeros as $\{a_1, \cdots, a_k\}$ in X_s, because s doesn't have any zeros. So if C is the zeroset of $a_1, \cdots, a_k$ in X, we indeed have $D = C \cap X_s$. □

Localization is important for two main reasons: it's easy to understand the relationship between A and A_s, and **the localizations form a basis for the topology on** X. (Here, a family of open subsets is a **basis** for a topology if every open set is a union of members of this family.) So any open set can be covered by localizations, but this doesn't necessarily mean all open sets are localizations!

Example 43

Let $X = \mathbb{A}^2_{x,y}$, and let $U = X - \{0, 0\}$. This is an open set, but it is not a localization.

To show this, we need to find an element s such that we can get U by inverting s. But if $s \neq 0$, inverting it doesn't make the point disappear, so we need to invert a function that is zero at the origin – every such polynomial vanishes on a curve.

How do we show that any open set U can be covered by localizations? If we want to show that $U = \bigcup X_{s_i}$, we should take complements: let $C = X - U$ be a closed set. Then the complement of any localization, $Z_i = X - X_{s_i}$, is the zeroset of s_i (because X_{s_i} is the set of points where s_i is equal to zero). So we know that

$$U = \bigcup X_{s_i} \iff C = \bigcap Z_i,$$

and by the definition of the Zariski topology, every closed set is the zeros of some elements in X, so we can indeed write U as a union of localizations.

We'll now move on to the concept of **morphisms**: let $X = \operatorname{Spec} A$ and $Y = \operatorname{Spec} B$, and say that $\phi : A \to B$ is a homomorphism. Then define the map $u : Y \to X$ (associated to the homomorphism ϕ) as follows: for any point $q \in Y$, we have a homomorphism $\pi_q : B \to \mathbb{C}$, which we can compose with ϕ to give a map $\pi_q\phi$. By definition of the spectrum of A, beause this is a homomorphism, it corresponds to some unique point $p \in X = \operatorname{Spec} A$: **define** $u(q) = p$.

What can we say about this to make it more clear what's going on? Let $\alpha \in A$ and $\phi(\alpha) = \beta \in B$. We want to evaluate α at a point p: by definition,

$$\alpha(p) = \pi_p(\alpha) = \pi_q\phi(\alpha) = \pi_q(\beta) = \beta(q).$$

So we define u such that evaluating α at p is the same as evaluating β at q.

Definition 44

A **morphism** $Y \to X$ is a map determined by the algebra homomorphism $\phi : A \to B$ as detailed above.

Let's think about this when we've chosen a specific presentation

$$A = \mathbb{C}[x]/(f), \quad x = (x_1, \cdots, x_n), \quad (f) = (f_1, \cdots, f_k).$$

Let τ be the canonical map from $\mathbb{C}[x]$ to A (modding out by the ideal (f)): composing this with ϕ gives us a map $\Phi : \mathbb{C}[x] \to B$. When can we tell that Φ comes from a map ϕ? Any homomorphism Φ comes from substituting in some values b_i for our x_is, and we need to make sure that everything in $\ker \tau$ is also in $\ker \Phi$ for the map ϕ to be well-defined (this is both necessary and sufficient). So this means that $\Phi(f_j) = 0$ for all j, meaning that $f_j(b) = 0$.

Proposition 45

In other words, constructing a homomorphism from $A = \mathbb{C}[x]/(f)$ to an arbitrary B, and thus constructing a morphism from $\operatorname{Spec} B$ to $\operatorname{Spec} A$, means we need to solve the equations $f_j(x) = 0$ for $x \in B$.

Example 46

Let $A = \mathbb{C}[x, y]/(y^2 - x^3)$, and let $B = \mathbb{C}[t]$. Then we can solve the equation $y^2 = x^3$ in B: one solution is $y = t^3, x = t^2$, and this defines a homomorphism $A \to B$ by doing exactly that (send (x, y) to (t^2, t^3)). And we know that $\operatorname{Spec} A$ is the cusp curve, and $\operatorname{Spec} B$ is the t-line, so we've found a morphism from the t-line into the cusp curve.

Example 47

Let $X = SL_2(\mathbb{C})$, and let $A = \mathbb{C}[a, b, c, d]/(ad - bc - 1)$. Then a map $A \to B$ means that we need to find a B-valued matrix with determinant 1. And in each case, we're finding a map into the special linear group.

We can also look at the blowup of the plane (which we studied a few lectures ago): let $A = \mathbb{C}[x, y]$, and let $B = \mathbb{C}[x, w]$. We map $A \to B$ by sending $(x, y) \to (x, xw)$. Then we again have a morphism: $\operatorname{Spec} B$ maps to $\operatorname{Spec} A$, and the w-axis is mapped to the origin, while the rest of the y-axis is not in the image.

10 February 24, 2020

Today, we'll talk about operations of a finite group on an affine variety, but let's start with a bit of review. Recall that if A and B are finite type domains, and $X = \operatorname{Spec} A, Y = \operatorname{Spec} B$, a homomorphism $\phi : A \to B$ gives us a **morphism** $u : Y \to X$ defined as follows: if $q \in Y$ is a point, this corresponds to a homomorphism $\pi_q : B \to \mathbb{C}$, which we can compose with ϕ to get a map $\pi_p = \pi_q\phi$. A way to think about this is that for any $\alpha \in A$, we can define

$$\alpha(p) = \pi_p(\alpha) = \pi_q(\phi\alpha) = [\phi\alpha](q).$$

So now let G be a finite group of automorphisms of B. An element $b \in B$ is **invariant** if $\sigma b = b$ for all $\sigma \in G$; then the set of invariant elements, $A = B^G$, is a subalgebra of B.

So for every automorphism $\sigma : B \to B$, we get a morphism $u_\sigma : Y \to Y$ (which we'll just denote σ). Then G operates on Y, but there's a bit of a problem: suppose we compose two automorphisms together, so we get the maps

$$B \xrightarrow{\sigma} B \xrightarrow{\tau} B.$$

But the corresponding maps backwards look like

$$Y \xleftarrow{\sigma} Y \xleftarrow{\tau} Y$$

(going from algebras to varieties reverses arrows), which yields a different composition. So an easy fix is as follows: let σ operate **on the left** on B, so we send b to σb, but let it operate **on the right** on Y, so q is sent to $q\sigma$. And this fixes the problem: a map $\tau\sigma$ sends $q \in Y$ to $q\tau\sigma$.

So suppose we have an element $\beta \in B$: what can we say about

$$\boxed{[\sigma\beta](q)}?$$

The definition of the regular function tells us that this is $\pi_q(\sigma\beta)$, and now this is a composition of functions $(\pi_q \circ \sigma)(\beta)$. And now note that if $p = q\sigma$, then this is equal to $\pi_p(\beta) = \beta(p)$, so this does give us $\boxed{\beta(q\sigma)}$. (So this is a way to move the σ back and forth.)

So if we map $A = B^G$ to B (for instance, via an inclusion), we have a map π from Y into $X = \operatorname{Spec} A$, which can be pretty interesting.

Example 48

Let $B = \mathbb{C}[y_1, \cdots, y_n]$, which means $Y = \mathbb{A}^n$, and let $G = S_n$ operate on the indices. Then B^G is generated by the elementary symmetric functions, so

$$A = B^G = \mathbb{C}[s_1, s_2, \cdots, s_n].$$

Then $X = \mathbb{A}^n_s$ is the affine space labeled by s.

If we have a point of X, which means we know the values of the symmetric functions $s_i = a_i$, then the y_i are the roots of the polynomial $x^n - a_1x^{n-1} + \cdots + (-1)^n a_n$. So the fibres of the map π from Y to X are the orbits of the group G operating on $\mathbb{A}^n_y$: the standard notation for the set of orbits is to use $X = Y/G$.

Example 49

Let $B = \mathbb{C}[y_1, y_2]$ and $Y = \operatorname{Spec} \mathbb{C}[y_1, y_2]$, and consider $G = \langle \sigma \rangle$ to be the cyclic group generated by the map

$$\sigma y_1 = \zeta y_1, \quad \sigma y_2 = \zeta^{-1} y_2,$$

where $\zeta = e^{2\pi i/n}$.

The only invariants of σ are those where the difference of the exponents is 0 mod n, so

$$A = B^G = \mathbb{C}[u_1, u_2, w]/(w^n - u_1 u_2),$$

where $u = y_1^n, u_2 = y_2^n, w = y_1 y_2$. (It's not hard to show that this generates all invariant polynomials.)

So suppose that we're given u_1, u_2, w which satisfy the relation, and say $u_1 \neq 0$. Then there are n possible values for y_1, and choosing this value fixes y_2 (because we know the value of $y_1 y_2$). So the n values of y_1 are the orbit of the original value under G, and thus $X = \operatorname{Spec} A$ is again equal to the set of the G-orbits. (The edge case is the origin, which is a fixed point under σ.)

Theorem 50

Let B be a finite-type domain, and let G be a finite group of automorphisms of B. Let $A = B^G$ be the invariant ring, and let $Y = \operatorname{Spec} B, X = \operatorname{Spec} A$. Then we have an inclusion $\pi : Y \to X$ (because A is a subset of B). Then

- A is a finite-type domain, and B is a finite A-module.
- The fibres of π are the G-orbits of π: and we have a bijective map $Y/G \to X$.

We'll do proof by example:

"Proof". Let $Y = SL_2(\mathbb{C})$ be the set of 2×2 complex-valued matrices with determinant 1. Then $B = \mathbb{C}[a, b, c, d]/(ad - bc = 1)$. Let $G = \langle \sigma \rangle$, where $\sigma^2 = 1$ and σ sends a point $P \in Y$ to its inverse P^{-1}. In other words,

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \mapsto \begin{bmatrix} d & -b \\ -c & a \end{bmatrix}.$$

These two matrices are identified in the orbits of G. To show the first point in the theorem, we need to write down some invariants: notice that $t = a + d, u = b^2, v = c^2, w = ad$ are all invariant under σ, and let $R = \mathbb{C}[t, w, u, v]/((w-1)^2 - uv)$. (because $ad - bc = 1$ in B, we also have the relation $(w-1)^2 = uv$ in R). Note that we have the chain of inclusions $R \subset A \subset B$.

Now a, d are roots of $x^2 - tx + w$, and b, c are roots of $x^2 - u, x^2 - v$ respectively. So let's write B down as a module over R: we have the elements $\boxed{1, a, b, c, ab, ac, bc, abc}$ (we don't need d because we get it from t and a), and any higher degree monomial in a, b, c will be divisible by at least one square, and we can use the equations to reduce the degree (for example, $a^2 - ta + w = 0$, so $a^2 = ta - w$). So everything in B can be written as a linear combination of these eight boxed elements, with coefficients in R.

Since R is a finitely-generated algebra (it's generated by t, w, u, v), it's Noetherian. Since A is an R-submodule of B (which is a finite R-module), A is a finite R-module (this is property of Noetherian rings). To prove that A is a finite type algebra, note that we have a finite number of generators of R as an algebra, and a finite number of generators of A as an R-module. So we have a finite number of generators as an algebra, and we've proved our first point.

Remark 51. *Note that R isn't the invariant ring here: there are other invariant polynomials, like bc, which are not in R.*

For the second point, we need to first (1) note that all points of Y in an orbit have the same image, then show that (2) points in different orbits have distinct images. Then we'll show that (3) every point in X is an image of a point of Y/G – that'll show that we have a bijection.

(3) is the most interesting: we will show that the map from Y to X is surjective. Let $p \in X$: then we have the maximal ideal $\mathfrak{m}_p$ of A at p. Because A is contained in B, we can look at $\mathfrak{m}_pB$, the "extended ideal," whose elements are combinations of the form

$$= \{z_1b_1 + \cdots + z_kb_k | z_i \in \mathfrak{m}_p, b_i \in B\}.$$

This is indeed an ideal of B, and every ideal except the unit ideal is contained in a maximal ideal. Say $\mathfrak{m}_pB \neq B$: then $\mathfrak{m}_pB$ is inside some maximal ideal $\mathfrak{m}_q$ in B corresponding to a point $q \in Y$.

So now we know that $\mathfrak{m}_pB \subset \mathfrak{m}_q$: consider $\mathfrak{m}_q \cap A$, which is an ideal of A (and not the unit ideal, because it doesn't contain 1). On the other hand, $\mathfrak{m}_p$ is contained inside $\mathfrak{m}_q \cap A$, and $\mathfrak{m}_p$ is a maximal ideal. So $\mathfrak{m}_p = \mathfrak{m}_q \cap A$, which means the map from Y to X indeed sends q to p, and we have a surjective map as desired.

So we just need to check whether $\mathfrak{m}_pB$ can be the unit ideal: can we write $1 = \sum z_ib_i$? The answer is no – suppose for the sake of contradiction, and sum over the orbit:

$$\sum_{\sigma \in G} \sigma(1) = \sum_{i,\sigma} \sigma(z_i)\sigma(b_i).$$

$\sigma(1) = 1$, and the z's are in the maximal ideal $\mathfrak{m}_p$, so $\sigma(z_i)$s are in A (the invariant ring). Therefore, this is

$$= \sum_{i,\sigma} z_i\sigma(b_i) = \sum_i z_i\alpha_i,$$

where $\alpha = \sum \sigma(b_i)$. So equating both sides, we have

$$|G| = \sum_i z_i\alpha_i,$$

where $\alpha_i \in A$ (because it's fixed under σ), $z_i \in \mathfrak{m}_p$. So $|G| \in \mathfrak{m}_p$, but $|G|$ is invertible in $\mathbb{C}$, which is a contradiction (because $\mathfrak{m}_p$ will contain 1). □

The central idea of this proof is to take advantage of certain invariants, which allow us to turn each generator of B into a finite set of generators of B as an R-module.

11 February 26, 2020

We'll talk for about three lectures about **projective varieties**. Recall that $\mathbb{P}^n$ is represented by the points $(x_0, \cdots, x_n) \sim (\lambda x_0, \cdots, \lambda x_n)$: then a Zariski closed set in $\mathbb{P}^n$ is the zeroset of a family of polynomials $f_1, \cdots, f_k$, where each f_i is **homogeneous** in $\mathbb{C}[x_0, \cdots, x_n]$. And now, just as in affine space, we have

$$V(f) = V(I), \quad I = (f_1, \cdots, f_k).$$

Lemma 52

An ideal I in $\mathbb{C}[x]$ is generated by homogeneous polynomials if and only if the homogeneous parts of f (of each degree) are in I.

Just like with $\mathbb{A}^n$, since the polynomial ring $\mathbb{C}[x]$ has the ascending chain condition on ideals, $\mathbb{P}^n$ has the descending chain condition on closed sets, so it is a Noetherian space. (In particular, this means that every closed set is a finite union of irreducible closed sets). Analogous to the affine case, an irreducible closed set is a **projective variety**.

There's just one thing to be careful about: the maximal ideal

$$M = (x_0, \cdots, x_n)$$

is called the **irrelevant ideal**, because its zeroset is empty – $(0, 0, \cdots, 0)$ is not a point in projective space.

Proposition 53

If $X \subset \mathbb{P}^n$ is a projective variety, and P is the ideal of polynomials

$$P = \{f | f = 0 \text{ on } X\},$$

then P is a prime ideal (different from M).

(The proof is analogous to the affine case.)

Example 54

A **hypersurface** is the zeroset of a single irreducible homogeneous polynomial $f(x)$ – this is a projective variety.

Example 55

The **Segre embedding** of $\mathbb{P}^m_x \times \mathbb{P}^n_y$ is a map $s : \mathbb{P}^m_x \times \mathbb{P}^n_y \to \mathbb{P}^N_w$, where we index our coordinates via

$$w = \{w_{ij}\}, 0 \le i \le m, 0 \le j \le n.$$

Then the point (x, y) maps to w via $w_{ij} = x_i y_j$.

The dimension of this image space is $N = (m+1)(n+1) - 1$ (the -1 comes from us being in projective space).

Proposition 56

The Segre map is injective, and the image is the zero locus of the **Segre equations**

$$w_{ij} w_{k\ell} = w_{i\ell} w_{kj}$$

for all indices i, j, k, ℓ.

Proof. First of all, these above equations obviously hold if $w_{ij} = x_i y_j$. Let w be a point in $\mathbb{P}^N$ which satisfies the Segre equations: then some $w_{ij} \neq 0$. We'll say it's w_{00} without loss of generality, and we'll set $w_{00} = 1$ by scaling.

This means x_0, y_0 are nonzero, and we'll set $x_0 = 1, y_0 = 1$. Then

$$w_{ij} w_{00} = w_{i0} w_{0j}$$

tells us that we should set $w_{ij} = x_i y_j$, and indeed this is the unique solution. □

Example 57

Let $m = n = 1$. Then we have a map $\mathbb{P}^1 \times \mathbb{P}^1 \to \mathbb{P}^3$ sending $(x_0, x_1) \times (y_0, y_1)$ to $(w_{00}, w_{01}, w_{10}, w_{11})$: then the image is the zero locus of $w_{00}w_{11} = w_{10}w_{01}$, which means $\mathbb{P}^1 \times \mathbb{P}^1$ is actually a quadric.

Example 58

The **Veronese embedding of $\mathbb{P}^n$** is a map $v : \mathbb{P}^n_x \to \mathbb{P}^N_v$, where we index coordinates via

$$v = \{v_{ij}\}, 0 \le i \le j \le n.$$

Then (x) maps to $v_{ij} = x_i x_j$.

Proposition 59

The Veronese map is injective, and its image is the zero locus of the equations

$$v_{ij}v_{k\ell} = v_{i\ell}v_{ij}, \quad 0 \le i \le k \le \ell \le j \le n.$$

Example 60

We can send $\mathbb{P}^1 \to \mathbb{P}^2$ under the Veronese embedding: (x_0, x_1) maps to the zero locus (x_{00}, x_{01}, x_{11}) where $v_{00}v_{11} = v_{01}^2$. (And this means that $\mathbb{P}^1$ is a conic.)

The Veronese embedding uses a quadratic basis, but we could take cubics or higher-degree monomials:

Example 61

Take the cubic Veronese embedding, which maps $\mathbb{P}^1 = (x_0, x_1)$ to $\mathbb{P}^3 = (z_{000}, z_{001}, z_{011}, z_{111}) = (x_0^3, x_0^2 x_1, x_0 x_1^2, x_1^3)$. We'll relabel the points in the image as (a, b, c, d): this forms a **twisted cubic** with the equations $b^2 = ac, ad = bc, c^2 = bd$.

It turns out that these are the 2×2 minors of the matrix $\begin{bmatrix} a & b & c \\ b & c & d \end{bmatrix}$: the points of the twisted cubic are often written as $(1, t, t^2, t^3)$ plus the point $(0, 0, 0, 1)$. And this is equivalent to $(u^3, u^2, u, 1)$ plus the point $(1, 0, 0, 0)$. where $u = t^{-1}$.

With this, let's go back to the Segre embedding: we know that $\mathbb{P}^m \times \mathbb{P}^n$ gets put into the zero locus of some homogeneous polynomial equations. **Is $\mathbb{P}^m \times \mathbb{P}^n$ a variety?**

We don't really have a definition for this yet, but we can look at $\mathbb{P}^m \times \mathbb{P}^n$'s image in the Segre embedding and see whether it's a variety in $\mathbb{P}^N$. But this is not obvious, and we need to ask the question "what's the Zariski topology in $\mathbb{P}^N$?"

Fact 62

If X and Y are varieties, then the Zariski topology on $X \times Y$ is **not** the product topology (unless one of X and Y is a point).

To show an example of this, consider $\mathbb{P}^1 \times \mathbb{P}^1$. The product topology is the **coarsest** topology (with fewest open sets) which makes the projections continuous. The closed sets in $\mathbb{P}^1$ are just finite sets, so **in the product topology**, the inverse image of a closed set in $\mathbb{P}^1 \times \mathbb{P}^1$ should be closed. This means the product topology just requires finite unions of "vertical/horizontal" lines and points to be closed, which is definitely not the Zariski topology.

Proposition 63

Let X and Y be irreducible topological spaces. Suppose there exists a topology on $\Pi = X \times Y$ such that

- The projections $\Pi \to X$, $\Pi \to Y$ are continuous. (This means that the topology on Π is at least as fine as the product topology.)
- For any $y \in Y$, the map from $\Pi_y = X \times y \to X$ is a homeomorphism, and similar for $\Pi_X \to Y$. (This means horizontal/vertical lines $X \times y$ and $Y \times x$ are "the same" as the original spaces X and Y.)

Then Π is an irreducible topological space.

(And this shows the result we want, because $\mathbb{P}^m$ and $\mathbb{P}^n$ are irreducible.)

Proof. Let C_1, C_2 be closed in Π such that $\Pi = C_1 \cup C_2$. We work with the complements (open sets) $W_1 = \Pi - C_1$ and $W_2 = \Pi - C_2$: it suffices to show that $W_1 \cap W_2 = \varnothing$.

Let U_1 be the image of W_1 in Y, and let U_2 be the image of W_2. Our first step is to show that if W is open in Π, then the image of U in Y is open. Intersect W_1 with a vertical fibre $\Pi_x = x \times Y$: then $W \cap \Pi_x$ is open in Π_x, and Π_x maps bijectively to Y. This means the image of U_x in Y is open, and now note that

$$W = \bigcup W_x \implies U = \bigcup U_x$$

must also be an open set. So now we know that U_1, U_2 are open – suppose that W_1 and W_2 are not both empty. Since Y is irreducible, $U_1 \cap U_2$ is not empty (or else the union of their complements in Y would be Y). Let $y \in U_1 \cap U_2$: U_1 is the image of W_1, so there is some point (x_1, y) in W_{1y} and (x_2, y) in W_{2y}. But W_{1y}, W_{2y} are open in Π_y, which maps bijectively to X. But that means Π_y is irreducible, which means $W_{1y} \cap W_{2y} \neq \varnothing$, and thus $W_1 \cap W_2 \neq \varnothing$, contradiction with the fact that $\Pi = C_1 \cup C_2$.

Thus either W_1 or W_2 is empty, which means $C_1 = \Pi$ or $C_2 = \Pi$, as desired. □

12 February 28, 2020

Today, we'll be talking about **morphisms** for projective varieties. We know that algebra homomorphisms give us affine varieties, but it's harder to make a definition for projective varieties (because points are equivalence classes in $\mathbb{P}^n$).

Example 64

Let's project a conic C onto a line $X = \mathbb{P}^1$. We choose our coordinates so that we project through the point $(1, 0, 0)$ onto the line $x_0 = 0$.

If we have a point $p = (a, b, c)$ on the conic, then πp will just be (b, c), because the line containing p and $(1, 0, 0)$ is $cx_1 = bx_2$. But what happens to the point q itself? We want to take a tangent line T at the point q, and see where that intersects X. To do more, we'll write out the explicit equation of the curve:

$$f = x_0x_1 + x_0x_2 + x_1x_2.$$

Then the tangent line is of the form

$$T : (\nabla f)_q \cdot x = 0,$$

where $f_0 = x_1 + x_2$, $f = x_0 + x_2$, $f_2 = x_0 + x_1$. The gradient evaluated at $q = (1, 0, 0)$ is $(0, 1, 1)$; thus, a point x is on the tangent line T if $x_1 + x_2 = 0$ and $x_0 = 0$, which gives the point $(0, 1, -1)$.

The central idea, though, is that there's no way to write down a formula in polynomials that works for **every point** on the conic! And there is no algebraic map from $\mathbb{P}^2$ to $\mathbb{P}^1$ except for constant maps, so there's something tricky to deal with here.

In general, when we look at projective varieties, **we'll call any nonempty open subset of a projective variety a variety as well**. (This is generally called a **quasiprojective variety**.) There are varieties that can't be put into a projective space, but we won't talk about them here.

Fact 65

An affine variety is quasiprojective, because we have the standard affine open sets $U^j = \{x_j \neq 0\} \subset \mathbb{P}^n_x$. And then $U^j \cong \mathbb{A}^n$ (because we need to normalize one of the coordinates to 1); another way to say that is that (defining some new variables)

$$U^j = \operatorname{Spec} \mathbb{C}\left[u_{0j}, \cdots, u_{nj}\right] = \operatorname{Spec} \mathbb{C}\left[\frac{x_0}{x_j}, \frac{x_1}{x_j}, \cdots, \frac{x_n}{x_j}\right].$$

Then an affine variety is a closed set in U^j, which gives us what we want.

Next, let X be a projective variety in $\mathbb{P}^n$, and let $X^j = X \cap U^j$ be a (closed) affine subset of U^j, ignoring the indices where $X \cap U_j = \varnothing$. (This won't happen if we pick generic coordinates.) Then $X_i \cap X_j \subset X_j$ is the subset of X_j where $x_i \neq 0$; that is, $u_{ij} = \frac{x_i}{x_j} \neq 0$.

We know that X^j is the Spec of some algebra A_j, which is a quotient ring of $\mathbb{C}[u_{0j}, \cdots, u_{nj}]$. We're doing some localizing here: note that

$$X^i = X^j = \operatorname{Spec} A_j[u_{ij}^{-1}] = \operatorname{Spec} A_i[u_{ji}^{-1}].$$

Thus, the field of fractions of A_j is the same for all j:

Definition 66

The **function field** K of X is the fraction field of any algebra A_j where $X_j \neq 0$. A **rational function** α is a nonzero element of K.

And open sets of this will still give the same function field: **all open subvarieties of a (projective) variety** have the same function field.

Definition 67

If α is a rational function on X, and $p \in X$, then α is **regular** at p if there is some index j and some $s \in A_j$ such that $\alpha \in (A_j)_s$ (adjoining s^{-1} to A_j).

We'll use these rational functions to define morphisms from a projective variety X to some projective space $\mathbb{P}^N$.

Definition 68

A **point with values in K** is the equivalence class of vectors $\alpha = (\alpha_0, \cdots, \alpha_N)$, with $\alpha_i \in K$ not all zero, where scaling the vector by any $\lambda \in K$ gives the same point.

If $p \in X$ is sent to some point with values in K, there are different cases that we should consider. If the α_i are all regular at p, and $\alpha_i(p)$ are not all zero, then define our map $\alpha : X \to \mathbb{P}^N$ via

$$\overline{\alpha} = (\alpha_0(p), \cdots, \alpha_n(p)).$$

On the other hand, if the α_i are not all regular or are all zero at p, then we look for a $\lambda \in K$ such that setting $\alpha_j' = \lambda a_j$ allows us to evaluate α' at p.

Example 69

Let's go back to our example from the beginning of class: what is the function field for the conic C?

Substituting our variables in, we find that (normalizing $u_0 = 1$)

$$u_1 u_2 + u_2 + u_1 = 0 \implies u_2 = -\frac{u_1}{u_1 + 1}.$$

This means that our function field is just

$$K = \mathbb{C}(u), u = u_1,$$

because we've already adjoined u_2 (a rational function of u_1) once we do u_1. And this gives us the map (just multiplying by various λs)

$$\pi(1, u_1, u_2) = (u_1, u_2) = \left(u, -\frac{u}{1+u}\right) = (1+u, -1)$$

which is a good function as long as our functions are regular ($u \neq -1$). Then writing $v = u^{-1}$, our map π sends a point $(1, u_1, u_2)$ to the point

$$(1 + v^{-1}, -1) \sim (v + 1, -v)$$

which is okay even at $v = 0 \implies u = \infty$.

Definition 70

A point α with values in K is a **good point** if for all $p \in X$, there exists a $\lambda \in K$ such that $\alpha_i' = \lambda a_i$ such that α_i' is regular at p for all i (and $\alpha_i'(p) \neq 0$ at at least one point). Then a **morphism** is a map defined by a good point.

This tells us how to define an **isomorphism** between two projective varieties X and Y: it is a morphism whose inverse is a morphism. (And remember that X and Y can be open subsets of projective varieties as well.)

Here's an example where this is useful:

Definition 71

An **affine open subset** of a variety X is an open subset U such that U is isomorphic to an affine variety $\operatorname{Spec} A$.

It's hard to know whether a subset is an affine open subset or not, especially in higher dimensions, but we have a nice result: let U, V be affine open subvarieties of a variety X (which sits inside some projective space and is not necessarily closed). Then $U \cap V$ is an affine open subvariety as well, and we'll prove this next time.

13 March 2, 2020

Before looking some more at affine open sets, we'll do a bit of review: a **variety** X is an open subset of a projective variety $\overline{X}$. The **function field** K of X is defined as follows: let $\overline{X}^i = \overline{X} \cap U^i$, where U^i is the standard open subset.

Then if $\overline{X}^i$ is not empty, then it is an affine variety: we can write $\overline{X}^i = \operatorname{Spec} A^i$, and then we define K to be the fraction field of A^i. But we don't really know what the affine open sets look like, so this is a bit clumsy.

Remark 72. *All open subvarieties of $\overline{X}$ have the same fraction field – note that we've been dealing with the closure this whole time.*

Nonzero elements of the function field are called **rational functions**: such a function f is **regular** at a point p if we can write $p \in \overline{X}^i$ for some i, meaning $\overline{X}^i = \operatorname{Spec} A^i$, and then the fraction f has a denominator that does not vanish at our point p – we then evaluate at our point.

A useful notion here is that of a **point of projective space with values in K**: this is just a collection $\alpha = (\alpha_0, \cdots, \alpha_r)$ with $\alpha_i \in K$ (not all zero) and the equivalence relation $\alpha \sim \lambda\alpha$ for any $\lambda \in K$. Then we can say that α is a point of $\overline{X}$ is it satisfies the defining equations of $\overline{X}$ – this helps us define a **morphism** $\underline{\alpha}$ from a variety X to projective space $\mathbb{P}^n$. Specifically, if we have a point $p \in X$ we look for a $\lambda \in K$ in the functional field such that

- $\alpha_i' = \lambda\alpha_i$ for all i,
- α_i' are not all zero, and they are all regular at p.

If this is possible for every point $p \in X$, α is called a **good point**, and then we just define the map

$$\underline{\alpha}(p) = (\alpha_0'(p), \cdots, \alpha_n'(p)).$$

$\underline{\alpha}$ gives us a morphism of X to $\mathbb{P}^n$: if the image is contained in subvariety Y of $\mathbb{P}^n$, then $\underline{\alpha}$ is a **morphism to Y**.

Lemma 73

Let $\{X^i\}$ be an open covering of X. Then if $\underline{\alpha}^i$ are morphisms from $X^i \to \mathbb{P}^n$, and $\underline{\alpha}^i$, $\underline{\alpha}^j$ are equal on $X^i \cap X^j$, then we can consistently define a morphism $\underline{\alpha}$ from X to $\mathbb{P}^n$.

Proof. The X^i have the same function field, and if $\underline{\alpha}^i$, $\underline{\alpha}^j$ are equal, then $\alpha^j = \lambda\alpha^i$ for some $\lambda \in K$. So we adjust α^j such that $\alpha^i = \alpha^j$ for all j (using different λs for each j, of course), and we can just let this function that we've defined be α. □

Definition 74

Let $X \subset \mathbb{P}^r, Y \subset \mathbb{P}^n$, and suppose we have a morphism $\underline{\alpha}$ from X to Y. Then $\underline{\alpha}$ is an **isomorphism** if it is bijective and its inverse is also a morphism.

We can use the above lemma to help us now:

Lemma 75

Suppose $\{Y_i\}$ is an open covering of Y. Let X^i be the inverse image of Y^i: then $\underline{\alpha}$ restricts to maps $\underline{\alpha}^i$ from X^i to Y^i. Then if $\underline{\alpha}^i$ is an isomorphism for all i, then α is an isomorphism.

Proof. Let $\underline{\beta}^i$ be the inverse morphism $Y^i \to X^i$. Then $\underline{\beta}^i = \underline{\beta}^j$ on $Y^i \cap Y^j$, because $\underline{\alpha}^i$ and $\underline{\alpha}^j$ are equal. This means β is a morphism as well by our previous lemma, meaning that $\underline{\alpha}$ is an isomorphism. □

So if we want to check for isomorphisms, this can be just done locally (on open subsets). Recall from last time than an open subset U of a variety X is an **affine open subset** if U is isomorphic to $\operatorname{Spec} A$ for some finite-type domain A.

Theorem 76

Let U, V be affine open subsets of a variety X. Then $U \cap V$ is an affine open subset.

Proof. Identify U with $\operatorname{Spec} A$ and V with $\operatorname{Spec} B$. To prove $U \cap V$ is an affine variety, we claim the coordinate algebra is $R = [A, B]$, the finite-type domain generated by A and B. (In other words, if $A = \mathbb{C}[a_1, \cdots, a_r]$ and $B = \mathbb{C}[b_1, \cdots, b_s]$, A and B both have fraction fields equal to K, the function field of X, and we have $\mathbb{C}[a_1, \cdots, a_r; b_1, \cdots, b_s]$.) We will show that if $W = \operatorname{Spec} R$, then $W \cong U \cap V$.

Since X is a variety, it sits inside some $\mathbb{P}^n$, and U, V sit inside X. A maps to R by inclusion, so we have a map $W \to U$ (and similarly we have a map $W \to V$).

The inclusion map $X \to \mathbb{P}^n$ is some morphism $\underline{\alpha}$, where α is a point of $\mathbb{P}^n$ with values in K (it defines the embedding). But then $\underline{\alpha}$ also gives us a morphism $W \to \mathbb{P}^n$: note that W can either map to U or V first, which means **the image of W under $\underline{\alpha}$ is contained in** $U \cap V$.

We want to show the map $\underline{\alpha}$ is an isomorphism. Choose an affine open set $Z \subset U \cap V$. We want Z to be a localization of U and of V: specifically, we wish to show that

$$Z = U_s = V_t,$$

where $U_s = \operatorname{Spec} A_s = \operatorname{Spec} A[s^{-1}]$ and $V_t = \operatorname{Spec} B[t^{-1}]$.

Lemma 77

Let $p \in X$. Then there exists an affine open subset Z contaning p in $U \cap V$ that is both a localization of U and V: $Z = U_s = V_t$.

Proof. We know that if $U = \operatorname{Spec} A$, then the localizations of U form a basis for the topology on U. So if $p \in U \cap V$, we can choose a nonzero $s \in A$ such that $p \in U_s \subset U \cap V$. Similarly, choose $t \in B$ so that $p \in V_t \subset U_s$ (we can do this because U_s is open). Finally, choose $r \in A$ so that $p \in U_r \subset V_t \subset U_s \subset U \cap V$.

Now $r \in V_t$, so we can write $r = bt^{-k}$ for some $b \in B$. We can now write the localization $V_{tb} = (V_t)_b = (V_t)_r$, and this last expression is just U_r (r is invertible in V_t). Thus, we've found the desired localization. □

Wwe can do this for every point $p \in X$, so we can cover $U \cap V$ with sets Z which are localizations of both U and V. Remember that $R = [A, B]$: then if $U_s = V_t$, we know that $A_s = B_t$, so

$$R_s = [A, B]_s = [A_s, B] = [B_t, B] = B_t = A_s.$$

This means our morphism $\underline{\alpha}$ yields an **isomorphism** from W_s to Z for any s, and now we can cover $U \cap V$ with these Zs. Therefore, $W \to U \cap V$ is an isomorphism (by consistency of the morphisms). □

We're a day behind schedule – we'll spend a day on the Grassmannian $G(2, 4)$, but we should read about the exterior algebra $\wedge V$ on our own, where V is a complex vector space.

14 March 4, 2020

We'll discuss the Grassmannian today:

Definition 78

The **Grassmannian** $G(k,n)$ parameterizes the subspaces U of dimension k in an n-dimensional vector space $V = \mathbb{C}^n$.

We should think of $G(k,n)$ as a set of points – one point for every subspace. For example, $G(1,n)$ is one-dimensional subspaces of $\mathbb{C}^n$, which is the same thing as $\mathbb{P}^{n-1}$. Similarly, $G(n-1,n)$ is the dual space $(\mathbb{P}^{n-1})^*$, which has the same dimension.

This means that the first interesting case not covered by the above examples is $G(2,4)$: these actually correspond to the lines in $\mathbb{P}^3$.

Remark 79. *Professor Mattuck at MIT did a lot of research on Grassmannians, even if he's known for a lot of teaching too.*

Let's try to describe this in another way. Suppose U is a 2-dimensional subspace of $V = \mathbb{C}^4$, and say that U has a basis u_1, u_2. This gives us a 2×4 matrix, and we can usually change the basis via row reduction to

$$\begin{bmatrix} 1 & 0 & * & * \\ 0 & 1 & * & * \end{bmatrix},$$

where the $*$ yields a 4-dimensional subspace. This is true whenever we can row reduce the matrix $\begin{bmatrix} u_1 \\ u_2 \end{bmatrix}$ in this way, which yields an open set W_{12} isomorphic to $\mathbb{A}^4$. If we can't row reduce with the first two columns, we can always do so with some other two columns (because the two rows are linearly independent). For example,

$$W_{24} = \left\{ \begin{bmatrix} * & 1 & * & 0 \\ * & 0 & * & 1 \end{bmatrix} \right\}.$$

There are six such open sets, and therefore we can write the Grassmannian $G(2,4)$ as the union of 6 affine spaces $\mathbb{A}^4$.

To say a bit more, say that (v_1, v_2, v_3, v_4) is a basis of $V = \mathbb{C}^4$. Then the **exterior algebra**

$$\wedge V = \wedge^0 V \oplus \wedge^1 V \oplus \cdots$$

is the (noncommutative but associative) algebra generated by V with the relations

$$vw = -wv \quad \forall v, w \in V.$$

(It's equivalent to just assume that $vv = 0$ for all $v \in V$: expand $(v+w)(v+w) = 0$.) Here, $\wedge^k V$ is the set of elements in $\wedge V$ where we multiply k elements of V together:

- $\wedge^0 V$ is just $\mathbb{C}$, so it has a basis of 1.
- $\wedge^1 V$ has dimension 4, and it's generated by v_1, v_2, v_3, v_4.
- $\wedge^2 V$ has dimension 6, and it's generated by $v_1v_2, v_1v_3, v_1v_4, v_2v_3, v_2v_4, v_3v_4$.
- $\wedge^3 V$ has dimension 4, and it's generated by $v_1v_2v_3, v_1v_2v_4, v_1v_3v_4, v_2v_3v_4$.
- $\wedge^4 V$ has dimension 1, and it's generated by $v_1v_2v_3v_4$.

In higher dimensions, there is nothing: if we try to multiply 5 things together, each monomial in v_1, v_2, v_3, v_4 will have at least two copies of one of the basis elements, so it just becomes 0. And it's easy to generalize this to an exterior algebra $\wedge V$ where V is an n-dimensional vector space.

Now note that

$$(v_1v_4)(v_2v_3) = -v_1v_2v_4v_3 = v_1v_2v_3v_4,$$

and similarly

$$(v_2v_3)(v_1v_4) = -v_2v_1v_3v_4 = v_1v_2v_3v_4.$$

So $\wedge^2 V$ is actually commutative – notice that this actually carries the structure of $G(2,4)$ inside it, up to scalars.

Definition 80

A **decomposable element** in $\wedge^2 V$ is an element w such that $w = uv$ for some $u, v \in V$.

Proposition 81

If we can write $w = \sum_{i<j} a_{ij}v_iv_j$, then w is decomposable if and only if

$$a_{12}a_{34} - a_{13}a_{24} + a_{14}a_{23} = 0.$$

Proof. We'll do the more interesting direction: if w satisfies the relation, then we can try

$$uv = (a_{12}v_2 + a_{13}v_3 + a_{14}v_4)(-a_{12}v_1 + a_{23}v_3 + a_{24}v_4).$$

(There's a negative sign because it's the negative of a_{21}.) And now we can just expand this to find

$$w = a_{12}^2v_1v_2 + a_{12}a_{23}v_2v_3 + a_{12}a_{24}v_2v_4 + a_{12}a_{13}v_1v_3 + 0 + a_{13}a_{24}v_3v_4 + a_{12}a_{14}v_1v_4 - a_{14}a_{23}v_3v_4 + 0$$

(remembering that $v_3v_3 = v_4v_4 = 0$). And then combining the two v_3v_4 terms with our relation, and canceling the a_{12} everywhere indeed gives us w. (And if $a_{12} = 0$, use a different collection of the indices.) □

Proposition 82

If a nonzero element $w \in \wedge^2 V$ can be written as u_1u_2, then (u_1, u_2) form a basis of a 2-dimensional subspace of V. (In other words, they are independent.)

(Otherwise, u_2 is a multiple of u_1, and $u_1u_1 = 0$.) So now we have a way to characterize our subspaces: a decomposable element of $\wedge^2 V$ is a subspace, which is a point in the Grassmannian. We just need to show that different w's give different subspaces:

Proposition 83

Let $w = u_1u_2, w' = u_1'u_2'$, and define U, U' to be the span of $(u_1, u_2), (u_1', u_2')$ respectively. Then if $U = U'$, w' is a scalar times w.

Proof. Since $U = U'$, we can write $u_1' = au_1 + bu_2$ and $u_2' = cu_1 + du_2$. Expanding,

$$u_1'u_2' = (ad - bc)u_1u_2 \implies w' = (ad - bc)w.$$

(the other terms are zeros because they are squares) □

In particular, this means that the Grassmannian is a collection of points that satisfy $a_{12}a_{34} - a_{13}a_{24} + a_{14}a_{23} = 0$, up to scaling. So $G(2,4)$ corresponds bijectively to a quadric in $\mathbb{P}^5$ with that above equation! And in particular, that identifies $G(2,4)$ as a variety. (Note that this means $G(2,4)$ has dimension 4, but that's something we already found out earlier.)

This helps us answer some other questions:

Problem 84

Let X be a surface in $\mathbb{P}^3$. (We'll use coordinates (x_1, x_2, x_3, x_4) here.) Does it contain a line ℓ?

Let's look at the case where the line is

$$\ell_0 : \{(x_1, x_2, 0, 0)\},$$

or equivalently $x_3 = x_4 = 0$. To find out if $\ell_0 \subset X$, say that X is the zero locus $f(x_1, x_2, x_3, x_4) = 0$, where f is an irreducible (homogeneous, because we're in projective space) polynomial of degree d. Then when we plug in x_3, x_4 into f, the polynomial must be identically zero (because x_1, x_2 can be anything, except $(0,0)$).

Well, plugging in $x_3 = x_4 = 0$ will in general give us the monomials

$$c_0 x_1^d + c_1 x_1^{d-1} + \cdots + c_d x_2^d,$$

where c_i are some of the coefficients of f. Therefore, for surfaces of a given degree, **containing a line is equivalent to setting $(d+1)$ coefficients equal to zero.**

What's another way to say this? The number of monomials in x_1, x_2, x_3, x_4 of degree d is $N = \binom{d+3}{3}$, so the set of surfaces S in $\mathbb{P}^3$ of degree d is parameterized by projective space of dimension $N-1$. (We want f to be irreducible, but this is an open subset of that projective space.) The lines that contain ℓ_0 form a linear subspace of S of **codimension** $d+1$, which means we have dimension $N - 1 - (d+1) = N - d - 2$.

But we care about containing **some** line, not just ℓ_0. However, the surfaces containing any line ℓ form a linear subspace of codimension $d+1$. Consider

$$\Gamma \subset G(2,4) \times S$$

where Γ corresponds to ordered pairs $([\ell], [X])$: ℓ is a line corresponding to a point in $G(2,4)$, and X is a surface, where $\ell \subset X$. The dimension of S is what we found earlier, and the dimension of $G(2,4)$ is 4, so we can add those dimensions together:

$$\dim \Gamma = 4 + (N - d - 2) = N - d + 2 = \boxed{\dim S - d + 3}.$$

- When $d = 1$, we have $\dim \Gamma = \dim S + 2$, so the fibres of the map $\Gamma \to S$ have dimension 2. Indeed, degree 1 hypersurfaces are planes, and they contain a 2-dimensional subspace of lines.
- When $d = 2$, we have $\dim \Gamma = \dim S + 1$, so the fibres have dimension 1. This is correct: quadrics are like hyperboloids, which have a one-dimensional family of rulings.
- When $d = 3$, we have $\dim \Gamma = \dim S$, so the fibres have dimension 0 most of the time. In other words, we just have finite sets – **most (generic) cubics have a finite set of lines**, and the right number is **27**. We'll try to be able to prove this! Remember that a generic quartic in $\mathbb{P}^2$ has 28 bitangents, and this is not an accident.

We don't always need to have a finite set of lines for all cubics – we can always take a cubic in three variables and draw all of the lines from that cubic to another point outside that plane, and then we'll have infinitely many lines.

15 March 6, 2020

We'll start talking about **dimension** today: there are two equivalent definitions, and we'll choose one of them.

Definition 85

Let X be a variety with function field K. The **dimension** of X, denoted $\dim X$, is the transcendence degree of K (that is, the maximum number of algebraically independent elements of K).

Another, more useful, definition is that $\dim X$ is the **longest possible length** k of a chain of closed subvarieties (irreducible closed subsets) in X of the form (notice there are $k+1$ subvarieties here)

$$X = C_0 > C_1 > C_2 > C_k = \text{point}.$$

Let's show that this definition is equivalent:

Proposition 86

All chains of closed subvarieties have length at most $\dim X$, and equality holds if we have a maximal chain.

We'll start by reviewing a concept: $\alpha_1, \cdots, \alpha_k \in K$ are **algebraically independent** if for all polynomials $f(x_1, \cdots, x_k)$, $f(\alpha) \neq 0$ – that is, we can't find a polynomial relation in the α_is. The **transcendence basis** is a set of elements of a maximal algebraically independent set (if finite). (The maximum number can be infinite, though it isn't when K is a function field).

All transcendence bases have the same order – this is fairly simple to prove by induction.

Lemma 87

Suppose $\alpha_1, \cdots, \alpha_k, \beta$ are elements of our function field K, and $\{\alpha_1, \cdots, \alpha_k\}$ are algebraically independent. Then adding β will keep our set algebraically independent if and only if β is not algebraic over $\mathbb{C}(\alpha_1, \cdots, \alpha_k)$.

Proposition 88

Let A be a domain (which is also an algebra), and let K be the fraction field of A. If the transcendence degree of K is finite, then there exists a basis of elements of A.

Proof. We proceed by induction – suppose $\alpha_1, \cdots, \alpha_k$ are elements of A that are algebraically independent but do not form a transcendence basis. Then there is some $\beta \in K$ which is not algebraic over the field $\mathbb{C}(\alpha_1, \cdots, \alpha_k)$. We can add β to our set (which we don't want to keep, because we want our basis to contain elements of A), which is of the form $\frac{b}{a}$ for $a, b \in A$. But a, b can't both be algebraic, or their quotient would be algebraic, so we can add either a or b to our transcendence basis. Therefore, we can always find another element in A to add, and we've created an algebraically independent set $(\alpha_1, \cdots, \alpha_{k+1})$. □

It's helpful to understand the dimensions along our chain of subvarieties $C_0 > C_1 > \cdots > C_k$:

Theorem 89 (Krull)

Let X be an affine variety of dimension n with $X = \operatorname{Spec} A$, and let $\alpha \in A$ be a nonzero element. Let V be the set of zeros of α in X – this forms a closed subset of X, which means it's a finite union of irreducible closed sets. Then every irreducible component of V has dimension $n-1$.

The proof is a bit harder than we might initially think:

Fact 90

This proof was presented during class, but it was simplified a bit too much and became incorrect. I've included the correct proof later in the notes.

Proof. Choose $\alpha_1, \cdots, \alpha_n \in A$ to be a transcendence basis for K. Specifically, we can assume α is not a scalar (otherwise V is empty, so the statement is vacuously true), so we can assume $\alpha = \alpha_n$.

Let W be an irreducible component of V. V is the union of finitely many components: let Z be the union of the others. "We want Z to be empty," so we localize: each of W and Z are defined by some equations, and in particular some elements of A are zero on Z. **Pick an $s \in A$ such that s is identically zero on Z but not identically zero on W**, and we'll look at $X_s = \operatorname{Spec} A[s^{-1}]$.

X_s has the same dimension as X, because it has the same field of fractions, and

$$X_s = X - (\text{zeros of } s).$$

In particular, $Z_s = Z \cap X_s$ is the empty set, but $W_s = W \cap X_s$ is not empty (because s is identically zero on Z, but not necessarily on W). Since X_s, W_s have the same dimensions as X, W respectively, we can replace them in our statement and use X_s, W_s instead.

Now since W (secretly $W \cap X_s$) is a closed subvariety of X, we can write $W = \operatorname{Spec} B$, where $B = A/P$ for some prime ideal P. In particular, the zeros of P in W are the zeros of α_n, which is exactly W. Now P and αA have the same zeros, so the Strong Nullstellensatz tells us that

$$\operatorname{rad}(\alpha A) = \operatorname{rad} P = P.$$

Remember that $\operatorname{rad} I$ is the set of elements a such that some power of a is in I. If $\beta \in A$ and β vanishes on W (meaning it's in the prime ideal P), then β^k is in αA for some k.

We want to show that $\dim W = \dim A - 1$: let $\overline{\alpha}_i$ be the residue of α_i in B (mod our prime ideal P), which gives us $(\overline{\alpha}_1, \cdots, \overline{\alpha}_{n-1})$. (Remember that $\overline{\alpha}_n = 0$ because α goes to 0 when we mod out by P.) We claim that these form a transcendence basis for $\overline{K}$, the fraction field of B – this would show that $\dim W = n - 1$.

There are two parts of this – we need to show they're algebraically independent, and also every other element is algebraic over $\mathbb{C}[\overline{\alpha}_1, \cdots, \overline{\alpha}_{n-1}]$.

First, we'll show that $\overline{\alpha}_i$s are algebraically independent. Say $f(x_1, \cdots, x_{n-1})$ is a polynomial with $f(\overline{\alpha}) = 0$: then $f(\alpha_1, \cdots, \alpha_{n-1}) \in P$, where P is the radical of αA. (? We'll postpone the rest of this proof for next time.)

On the other hand, let's show that $\overline{\alpha}_1, \cdots, \overline{\alpha}_{n-1}, \overline{\beta}$ are always algebraically dependent. Suppose otherwise: then $\alpha_1, \cdots, \alpha_{n-1}, \beta$ form a transcendence basis in K if β is not algebraic over $\mathbb{C}(\alpha_1, \cdots, \alpha_{n-1})$. We know that β must be algebraic over $\mathbb{C}(\alpha_1, \cdots, \alpha_n)$, so we can write

$$\sum c_i(\alpha)\beta^i = 0,$$

where $c_i(\alpha)$ are rational functions of α. Clearing denominators, we can assume they are polynomials, so we have some polynomial $f(x_1, \cdots, x_n, y)$ such that $f(\alpha_1, \cdots, \alpha_n, \beta) = 0$. This tells us that (modding out by P)

$$f(\overline{\alpha}_1, \cdots, \overline{\alpha}_{n-1}, 0, \overline{\beta}) = 0,$$

We can assume that x_n does not divide f, so that when we plug in 0, f is not identically zero. But now we have a polynomial relation between $\overline{\beta}, \overline{\alpha}_1, \cdots, \overline{\alpha}_{n-1}$, which is a contradiction. □

16 March 9, 2020

We'll have a quiz next Wednesday during class – we'll make the next problem set due after spring break, as long as MIT isn't canceled by then.

By the way, the proof of Krull's theorem last time was incorrect, because we simplified the proof too much and it became incorrect. So we'll leave it for now and come back to it later on.

Today, we'll start with the Nakayama Lemma. Let A be a ring and $P = (p_{ij})$ be an $n \times n$ matrix with entries in A. Then there exists another A-matrix C such that $CP = (\det P) \cdot I$. Here, C is the **cofactor matrix**: the entry C_{ij} is $(-1)^{i+j}$ times the determinant of P when we remove the jth row and the ith column. What's important is that this formula works in any ring. As an example, if $P = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$, then we have $C = \begin{bmatrix} d & -b \\ -c & a \end{bmatrix}$ (we should remember to take the transpose after looking at the matrix minors).

So now let M be an A-module, and let $v = (v_1, \cdots, v_n)$ be a vector with entries in M. Let P be an $n \times n$ A-matrix: then v is an **eigenvector** with value λ if

$$Pv = \lambda v \implies (\lambda I - P)v = 0.$$

Now if we let C be the cofactor matrix for $(\lambda I - P)$, we know that $C(\lambda I - P)v = \det(\lambda I - P)v = 0$. In other words, we know that $p(\lambda)$, the **characteristic polynomial** $\det(\lambda I - P)$, kills the vector v.

Proposition 91 (Nakayama Lemma)
Let M be a finite A-module, and let J be an ideal of A. If $JM = M$, then there exists $z \in J$ such that $(1-z)M = 0$ – in other words, $zm = m$ for all $m \in M$.

Proof. Let $v_1, \cdots, v_n$ generate the module M, so every element of M is a combination of the v_is with coefficients in A. Let $v = \begin{bmatrix} v_1 \\ v_2 \\ \vdots \\ v_n \end{bmatrix}$. Since $JM = M$, each generator v_i can be written as a sum of terms jm, where $j \in J$ and $m \in M$, and now each m is an A-linear combination of the v_js. Moving the coefficients from the v_js to the j's (which is allowed by closure of the ideal J), we find that

$$v_i = \sum_j p_{ij} v_j,$$

where the coefficients p_{ij} are all in J. Thus v is an eigenvector of the matrix $P = (p_{ij})$ with eigenvalue 1.

This means that if we have our characteristic polynomial $p(t) = \det(tI - P)$, then $p(1)v = 0$. But every element in M is generated by the v_is, so $p(1)m = 0$ for all $m \in M$. Now we just need to find an element z: set

$$\det(I - P) = 1 - z,$$

and notice that the left hand side is some polynomial z in the p_{ij}s (which are all in J), so $z \in J$, as desired.

For example,

$$\begin{bmatrix} 1-a & b \\ c & 1-d \end{bmatrix} = (1-a)(1-d) - bc,$$

so this gives the element $z = a + d + bc - ad$. □

Corollary 92

Let A be a Noetherian domain, and let I, J be ideals of A. If $IJ = I$, then either I is the zero ideal or J is the unit ideal.

Proof. Since A is Noetherian, I is a finitely generated ideal, which means I is a finite A-module. Therefore, the Nakayama Lemma tells us that there is a $z \in J$ such that $(1-z)I = 0$. Therefore, either $I = 0$ or there is an element in I. In the latter case, we must have $z = 1$ (because there are no zero divisors in a domain), meaning J is the unit ideal because $z = 1 \in J$. □

Corollary 93

Let A be a Noetherian domain, and let $x \in A$ be **not a unit**. Then

$$\bigcap x^n A = (0).$$

In other words, if $y \neq 0$ is in A, then x^n does not divide y for sufficiently large n.

Proof. If $x = 0$, the statement is clearly true. Otherwise, let $I = \bigcap x^n A$, and let $J = xA$. Then $JI = I$ (we just multiply an extra x in every term), and J is not the unit ideal, so I is the zero ideal by the previous corollary. □

This turns out to be true when our ideal J is not just a principal ideal, but it's hard to show in that case. Also, note that the assumption that A is a domain is important: a counterexample otherwise is to take $A = \mathbb{C} \times \mathbb{C}$ and $x = (1, 0)$.

We'll look at something a bit different now:

Definition 94

Let $A \subset B$ be domains. Then an element $\beta \in B$ is **integral** over A if it is the root of a monic polynomial $f(x) = x^n + a_1 x^{n-1} + \cdots + a_n$, where $a_i \in A$.

The set of elements in B that are integral over A form a **subring of** B, which we can check by verifying all of the closure assumptions.

Lemma 95

Let A be a Noetherian ring. Then β is integral if and only if $A[\beta]$ is a finite A-module.

Proof. If β is integral, we have the relation $\beta^n =$(polynomial in β), so $(1, \beta, \beta^2, \cdots, \beta^{n-1})$ generate $A[\beta]$.

For the other direction, if $A[\beta]$ is a finite A-module, then we keep picking powers of β until they generate $A[\beta]$. Then we have some finite list of powers of β, so we can write β^N in terms of them, giving us a monic polynomial. □

Lemma 96

Let $A \subset B$ be finite type domains. Then B be a finite A-module if and only if all elements of B are integral over A.

Proof. Any submodule of B is a finite A-module, so all the elements are integral (because $A[\beta]$ is always a submodule of B). For the other direction, we know that B is generated by finitely many elements as an algebra. If each one is integral, then B is a finite module because we adjoin one of them and it's a finite module, then add the next one, and so on. □

Example 97

Let $A = \mathbb{C}[x]$, and let $B = A[y]/(f(y))$. Here, f is of the form

$$f(y) = a_0 y^n + a_1 y^{n-1} + \cdots + a_n,$$

where $a_i \in A$ are functions of x.

If $a_0 = 1$ (or any other nonzero complex number), then the residue of y in B is integral. Thus, every element of B is integral over A. But if a_0 is not a constant, it will vanish at some points where $a_0(x) = 0$. The fibre over any x_0 is a finite set of points, but at each of the points where a_0 vanishes, the degree of f goes down. So that corresponds to a vertical asymptote, where y goes to $\pm\infty$ as $x \to x_0$. We don't like these vertical asymptotes, which explains why we care about the idea of β being integral!

For example, when $f(x, y) = xy - 1$, then $a_0(x) = x$, and indeed as $x \to 0$, $y \to \infty$. More generally, if we write

$$f = a_0 y^n + \cdots + a_n,$$

we have

$$\frac{f}{a_0} = g = y^n + \frac{a_1}{a_0} y^{n-1} + \cdots + \frac{a_n}{a_0}.$$

If $a_0 \neq 0$, then this just looks like f, but when $a_0 = 0$, at least one of the coefficients goes to ∞. Since the coefficients are symmetric functions in the roots, this means at least one root must go to infinity.

Theorem 98 (Noether Normalization Theorem)

Let A be a finite-type domain. Then there exist $y_1, \cdots, y_r$ in A that are algebraically independent, such that A is a finite-module over a polynomial subring $\mathbb{C}[y_1, \cdots, y_n] = P$ (so all the elements of A are integral over the polynomials).

Proof. Choose generators $x_1, \cdots, x_n$ for our algebra A. If they are algebraically independent, then we're done – P is equal to the polynomial ring. Otherwise, let f be the polynomial that evaluates to 0 for our variables x_i. We'll use the following result:

Lemma 99

There exists a linear change of variables so that $f(x_1, \cdots, x_n)$ becomes a monic polynomial in x_n of deg $d = \deg f$.

Proof of lemma. Let $x_i = x_i' + c_i x_n'$, where the c_is are to be determined. (This is indeed invertible.) Take the highest-degree homogeneous part of f: then

$$f(x_1, \cdots, x_n) = f(x_1' + c_1 x_n', \cdots, x_n' + c_n x_n').$$

Our goal is to show that f can be monic in x_n'. The coefficient of degree d comes from putting in $x_1' = x_2' = \cdots + x_{n-1}' = 0$ and $x_n' = 1$: we'll get

$$f(c_1, \cdots, c_{n-1}, 1 + c_n),$$

which can be nonzero if we choose the c_is appropriately, and then we can normalize to finish. □

Then the Normalization Theorem becomes very simple: make that change of variables, and let R be the algebra $\mathbb{C}[x_1, \cdots, x_{n-1}] \subset A$. But we chose a polynomial f earlier that evaluates to zero at our point $(x_1, \cdots, x_n)$, and x is monic, so x_n is integral over R. Therefore, A is a finite R-module.

And now we induct: we may assume R is a finite module over the polynomial subring P. Then A must be a finite module over P, as desired. □

This turns out to be extremely useful, and we'll see applications shortly.

17 March 11, 2020

There won't be a quiz next week. Professor Artin will post a problem set for us to work on for the next few weeks. Lectures after this week might be done by Zoom, but things are a bit unclear for now. (We'll still have lecture on Friday.)

We'll jump ahead a bit for today's lecture, because it'll be hard to get things organized. First, a bit of review: if $A \subset B$, we say that B is **integral over** A if every element $b \in B$ satisfies a monic polynomial relation with coefficients in A. In particular, if A and B are finite-type domains, then B is integral over A if and only if B is a finite A-module.

It's also possible to do this when B is an A-module, but then we're not necessarily working with an inclusion anymore:

Definition 100

Let $Y = \operatorname{Spec} B$ and $X = \operatorname{Spec} A$. If there is a homomorphism $\phi : A \to B$, then we call the morphism $u : Y \to X$ a **finite morphism**. If $A \subset B$ and B is a finite A-module, then $Y \to X$ is an **integral morphism**.

Example 101

Let $Y = \operatorname{Spec} B$ be a closed subvariety of $X = \operatorname{Spec} A$. Then the inclusion map $Y \to X$ is a finite morphism, because the corresponding homomorphism $\phi : A \to B$ is a quotient map, and B is clearly a finite A-module because it is generated by $\{1\}$. (Here, it's helpful to note that B is an A-module in the sense that it is closed under multiplication by $\phi(A)$.)

We can also extend our definitions of finite and integral morphisms beyond the affine case as well. Any morphism $u : Y \to X$ is a **finite (resp: integral) morphism** if we can cover X by affine open sets, and for every such affine open set $X' = \operatorname{Spec} A \in X$, the inverse image $Y' = \operatorname{Spec} B$ is affine, and $u : Y' \to X'$ is **finite (resp: integral)**. In general, this means we only need to check on the affine open sets! To show that this definition is consistent with the previous one, we need to look at the localizations (which form a basis for the topology), though the proof is a bit long because we need to check a lot of different conditions.

For now, we'll move on:

Theorem 102 (Chevalley's finiteness theorem, version 1)

Let Y be a closed subvariety of $\mathbb{P}^n \times X$. If the projection $\pi : Y \to X$ has finite fibres, then π is a finite morphism.

Theorem 103 (Chevalley's finiteness theorem, version 2)
Let X, Y be projective varieties, and say that $u : Y \to X$ is a morphism. Then if the fibres of u are finite, then u is a finite morphism.

Let's first show that version 1 implies version 2. Let Γ be the **graph** of u, which consists of points of the form $(y, u(y))$. Then we have $Y \subset \mathbb{P}^n$, and $\Gamma \subset Y \times X \subset \mathbb{P}^n \times X$. Since Γ is a closed subvariety of $\mathbb{P}^n \times X$ and Γ is isomorphic to Y, we can compose the isomorphism with the morphism from Γ to X.

Proof of Chevalley (version 1). We'll prove this in the case of $\mathbb{P}^1$ – since $\mathbb{P}^n$ just has more indices, we'll avoid those. We know that Y sits inside $\mathbb{P} \times X$, and π is the projection of a point down to X. The graph of Γ looks like some curve in the X-axis and $\mathbb{P}^1$-axis: pick coordinates (y_0, y_1) in $\mathbb{P}^1$. Let X^∞ be the points at infinity – that is, where $y_0 = 0$.

In case 1 (this is the main case), we have that $Y \cap X^\infty = \varnothing$. Because we don't hit the line at infinity, we can assume X is affine, so $X = \operatorname{Spec} A$ for some A.

Consider the two standard open sets $U^0 \times X$ (where $y_0 \neq 0$) and $U^1 \times X$ (where $y_1 \neq 0$). Then we can say that

$$U^0 \times X = \operatorname{Spec} A[u], \quad U^1 \times X = \operatorname{Spec} A[v],$$

where $u = \frac{y_1}{y_0}, v = \frac{y_0}{y_1} = u^{-1}$. In this case, because $Y \cap X^\infty = \varnothing$, we know that

$$Y^0 = (U^0 \times X) \cap Y = Y.$$

The secret is to look at $U^1 \times X$ first: this can be written as $\operatorname{Spec} R_1$, where $R_1 = A[v]$. Then $Y^1 = (U^1 \times X) \cap Y$ is a closed subset of $U^1 \times X$, and it is irreducible (because it is an open subset of Y, which is irreducible), so Y^1 is the locus of some prime ideal Q. (Remember that when we work in $U^1 \times X$, we're leaving out the parts of the space where $y_1 = 0$.)

Now X^∞ is the locus where $v = 0$, and $Y^1 \cap X^\infty$ is empty because $Y \cap X^\infty$ is empty. But we have two closed subsets that don't intersect, which must generate the unit ideal. Since X^∞ is generated by v, we have that

$$Q + vR_1 = (1) = R_1.$$

Then there exist $f(v) \in Q, g(v) \in R_1$ in $R_1 = A[v]$ (we're writing everything as polynomials in v, with coefficients in A) such that

$$f(v) + vg(v) = 1.$$

Write out $f(v) = a_0 v^n + \cdots + a_n$; we know that $a_n = 1$ by matching constant coefficients. But now we can write everything in terms of $\frac{1}{v} = u$ instead: multiplying this through by a high power of u to clear all denominators, and we'll find that the left hand side is a **monic** polynomial in u, because the original constant term was 1.

To understand where this polynomial f vanishes, note that if we leave out both 0 and ∞ in the projective coordinate, we can go back and forth between u and v. This monic polynomial vanishes on all of $Y - X^\infty$, and thus it vanishes on Y by closure. That means that u is integral over A, but we can also write B as a quotient of $A[u]$ because Y is a closed affine subset of $U^0 \times X$. Thus B is a finite A-module, and we've proved the desired result.

In case 2, Y does intersect X^∞. Then the plan is to cover X by localizations so that we can get away from our problem: for any $p \in X$, $p \times \mathbb{P}^1$ intersects Y at only finitely many points. Since the localizations form a basis, we can cover X by $X^i = X_{s_i}$ so that after a change of coordinates, we localize the points $p \in X$ where Y intersects X^∞ (this is a closed set). Then we have a **finite number of localizations** (because of the finite fibre assumption)

$$Y^i = (\mathbb{P}^i \times X^i) \cap Y,$$

and this doesn't meet X^∞. Then $X^i = \operatorname{Spec} A_i$, and Y^i is the inverse image of X^i under π, so $Y^i = \operatorname{Spec} B_i$, where B_i is a finite A_i-module from case 1.

And the localizations cover X, so $s_1, \cdots, s_r$ generate the unit ideal of A.

Lemma 104

With all of the notation above, we have that $Y = \operatorname{Spec} B$ and B is a finite A-module.

This kind of argument often comes up in algebraic geometry, so we should pay attention to it here.

Proof of lemma. Let $R = \bigcap B_i$ – all of the B_is have the same fraction field, so we can take the intersection there. We'll show that $Y = \operatorname{Spec} R$ and R is a finite A-module.

Localize R, so

$$R_{s_j} = \left(\bigcap (B_i)\right)_{s_j}.$$

If we invert s_j in all of the B_is separately, we get something at least as big, so

$$R_{s_j} \subset \bigcap ((B_i)_{s_j}).$$

Thus, for all $\beta \in B$ in the right-hand set, we can write

$$\beta = b_i s_j^{-k}, \quad b_i \in B_i,$$

where we use the same k for all i for simplicity. Thus $\beta s_j^k = b_i$ for all i, and therefore βs_j^k is in the intersection of the B_is, which is R. Therefore, every element $\beta \in \bigcap ((B_i)_{s_j})$ is also in R_{s_j}, so we have

$$R_{s_j} = \bigcap ((B_i)_{s_j}).$$

B_i is the coordinate algebra of Y^i, so

$$Y^i \cap Y^j = \operatorname{Spec} B^i \cap \operatorname{Spec} B^j = \operatorname{Spec} (B_i)_{s_j},$$

because we need a function that vanishes on the complement of Y^j. And notice that $(B_j)_{s_j} = B_j$ (s_j is already invertible in B_j), which means that

$$Y^i \cap Y^j = \operatorname{Spec} (B_i)_{s_j} = \operatorname{Spec} (B_j)_{s_i}.$$

Thus $(B_i)_{s_j} = (B_j)_{s_i}$, and both of these are at least as big as B_j. But when $i = j$, we get equality, so indeed we have

$$\bigcap (B_i)_{s_j} = B_j.$$

So R is a nice ring, because **localizing gives us the B_js**. Let $\tilde{Y} = \operatorname{Spec} R$: we wish to show that $\tilde{Y} = Y$. When we localize $\tilde{Y}_{s_j}$, we end up with $\operatorname{Spec} R_{s_j} = \operatorname{Spec} B_j = Y^j$. Thus $\tilde{Y}_{s_j}$ is isomorphic to Y^j, and the former is a covering of $\tilde{Y}$ while the latter is a covering of Y. Thus we've shown that Y and $\tilde{Y}$ are indeed isomorphic.

To finish, note that $B_i = R_{s_i}$ is a finite A_i-module by construction for each i, so we can pick finitely many generators for each R_{s_i}. Then we can multiply by s_i (which is invertible in A_i), so that our generators are all elements of R, and we do this for all (finitely many) localizations i. Putting all of these together yields a set of generators for R. □

And now B is again a finite A-module, completing the proof.

□

18 March 30, 2020

Fact 105

Classes from this day onward were conducted over Zoom.

We're going to talk about **double planes** today, specifically projective double planes (we can read about affine double planes on our own). Throughout this class, we'll let $X = \mathbb{P}^2$ with coordinates x_0, x_1, x_2, and we'll let y be another variable of **weight** d: in other words, we'll let $f(x)$ be a homogeneous polynomial of degree $2d$, and we say that $y^2 = f(x)$. Then we'll look at the zero locus of $y^2 - f(x)$ in the **weighted projective space**, which is defined by the relation

$$(x_0, x_1, x_2, y) = (\lambda x_0, \lambda x_1, \lambda x_2, \lambda^d y).$$

Indeed, $y^2 - f(x)$ can be thought as a homogeneous degree $2d$ polynomial if y has weight d:

Definition 106

A **double plane** is a zero locus Y for $y^2 - f(x)$ in the weighted projective space.

Since X is the ordinary projective plane, we have the projection map $\pi : Y \to X$ sending (x_0, x_1, x_2, y) to (x_0, x_1, x_2). The fibres of this map generally have two points (x, y) and $(x, -y)$, but these are the same points when $f(x) = 0$. So the **branch locus** Δ corresponds to the points in $\mathbb{P}^2$ where the fibres of π are just a single point $(x, 0)$. And there's an automorphism $\sigma : Y \to Y$ of the double plane given by

$$\sigma(x, y) = (x, -y);$$

this is an order 2 automorphism.

Since our map π is two-to-one, the image of Y will have dimension 1 as a subvariety of X. Let C be a curve in X which is the zero locus $h(x) = 0$ of an irreducible homogeneous polynomial h. There are three cases here:

- C **ramifies**, which means that C is contained in the branch locus. This means that h divides f – this only occurs for a finite number of curves. In this case, $\pi^{-1}C$ will map bijectively to C, because every point in C corresponds to one point in $\pi^{-1}(C)$. (If we were talking about schemes, we'd want to say that the fibre is "two times the curve," because $y^2 = 0$. We don't have to worry about this, though.)
- C **splits**, which means that $\pi^{-1}C$ is made up of two curves $D_1 \cup D_2$ in Y. Applying σ to Y will switch the curves D_1 and D_2 with each other, and $D_i \to C$ are generically bijective for both $i = 1$ and $i = 2$.
- C **doesn't split**, which means that $\pi^{-1}C$ is an irreducible curve. This happens when the ideal

  $$(y^2 - f, h)$$

  is a prime ideal in $\mathbb{C}[x, y]$.

All of this is analogous to what happens when we adjoin a quadratic element δ to $\mathbb{Z}$ (this is material from 18.702):

Example 107

Consider $\mathbb{Z} \subset \mathbb{Z}[\delta]$, where $\delta^2 = -5$.

There are a few different possibilities here, which correspond to the above cases:

- Notice that the equation $y^2 = -5$ becomes $y^2 = 0$ when we're looking mod 5, so that explains why the prime $p = 5$ ramifies in $\mathbb{Z}[\delta]$.
- On the other hand, $p = 3$ splits as
$$(3) = (3, 1 + \delta), (3, 1 - \delta).$$
Since (3) does not split as a principal ideal, it makes it harder to analyze what the ideals actually look like in $\mathbb{Z}[\delta]$. So it's generally not very easy to say what the ideals look like in the case above where C splits. Meanwhile, $p = 41$ does split into principal ideals:
$$(41) = (6 + \delta)(6 - \delta).$$
- Finally, the prime $p = 7$ remains prime, which corresponds to the third case above.

So if we go back to the double plane $y^2 = f(x)$, remember that $f(x)$ has even degree. Let's take the particular example where $d = 2$ and $\deg f = 4$: we want to ask **when a line $C = L$ splits**.

We can choose our coordinates so that the line is $\{x_0 = 0\}$. Since L splits, we can say that the ideal $(y^2 - f, x_0)$ is not prime. Taking this ideal mod x_0, we can define $\overline{f}(x_1, x_2) = f(0, x_1, x_2)$, and then we can rephrase our statement as saying that $(y^2 - \overline{f})$ is **not a prime ideal** in $\mathbb{C}[x_1, x_2, y]$. Since $\mathbb{C}[x_1, x_2, y]$ is a UFD, this is true if and only if $y^2 - \overline{f}$ is **not irreducible**.

Since $y^2 - \overline{f}$ is quadratic in y, a factorization will look like

$$y^2 - f = (y + \alpha)(y - \alpha) \implies \overline{f} = \alpha^2$$

(each factor has to be linear in y, and the polynomial is monic in y). So the line L splits if and only if $\overline{f}$ is a square. This means that all of the roots of f have to have order 2, so Δ has to be **tangent** to the line L. So in this special case that we're talking about with $d = 2$, **lines that split correspond to bitangents of the branch locus** Δ, so that both of the two zeros have order 2. (It could also be a single fourfold tangent, but that's not generic, so we'll ignore it.) Note similarly that a line L that splits in a double plane with $d = 3$ is a tritangent for the same reason. But generic curves don't have tritangents, because there are finitely many bitangents!

With this, we can explain the following famous property:

Proposition 108

Every smooth cubic surface contains 27 lines.

Proof. Suppose we have a cubic surface $S \subset \mathbb{P}^3$, and we label $\mathbb{P}^3$ with coordinates (x_0, x_1, x_2, z) (it's not weighted, but we do want to distinguish one coordinate). Choose coordinates so that $q = (0, 0, 0, 1)$ **is a point on the curve** S: then the defining equation of S must be of the form

$$az^2 + bz + c = 0,$$

where a, b, c are homogeneous polynomials of degree $1, 2, 3$ in the x-variables. (Because $(0, 0, 0, 1)$ is on the curve, the coefficient of z^3 is zero.)

Now the discriminant of $f = b^2 - 4ac$ is a homogeneous quartic, and thus we can describe the zero locus similarly to the locus of the quartic double plane

$$Y : y^2 = f,$$

where the quadratic formula tells us that

$$z = \frac{-b+\sqrt{f}}{2a} = \frac{-b+z}{2a} \implies y = 2az + b.$$

(The zero loci here are **birational**.) So the correspondence between S and Y makes sense whenever $a \neq 0$ – let's look at the $a = 0$ **special case** first. Then we can find a line L_0 in X which splits in Y, because the defining equation of S looks like

$$\overline{b}z + \overline{c} = 0$$

modulo a. The locus of $a = 0$ then defines a plane H_0 in $\mathbb{P}^3$ with coordinates (x_0, x_1, x_2, z), and this plane contains the point q. Meanwhile, the locus of $\overline{b}z + \overline{c} = 0$ is a (non-generic) cubic curve in H_0 – this equation is singular at the point q, because there is no z^3 coefficient. (To verify this, we can just suppose that a is the coordinate x_0 by a change of coordinates, and we can note that it's not possible for that cubic curve to be the union of a line and a quadratic if S is generic.

But now let's look at the case where $a \neq 0$, we can take L to be another bitangent to the curve $f = b^2 - 4ac = 0$. L now intersects our first bitangent L_0 in one point, and the map between S and Y is defined except at a single point of L. Since L splits in Y (it's a bitangent), we know that it is the union of a line and a quadratic. Since there are 27 bitangents distinct from L_0, this means that there are 27 lines on a cubic, as desired. □

19 April 1, 2020

There's a lot on the syllabus for today – we'll talk about valuations and then discuss a little about normalization.

Definition 109

Let K be a field. A (discrete rank 1) **valuation** on K is a surjective homomorphism from $K^\times \to \mathbb{Z}^+$, such that $v(\alpha\beta) = v(\alpha) + v(\beta)$ for all $\alpha, \beta \in K$, and $v(\alpha + \beta) \geq \min(v(\alpha), v(\beta))$.

Example 110

Let $K = \mathbb{C}(t)$, and let $t = a$ be a point in A^1_t.

We can define the valuation v_p with the following process: if $f(t)$ is a polynomial, and $t - a$ divides f k times, then $v_p(t) = k$. And if we have a rational function $\alpha = \frac{f}{g}$, we'll just define $v_p(x) = v_p(f) - v_p(g)$.

It turns out that all valuations of $K = \mathbb{C}(t)$ are either v_p for some p or the valuation at the point at ∞ (which we get by working with $\frac{1}{t}$ instead).

Definition 111

The **valuation ring** R of a valuation v is the set of $\alpha \in K$ such that $v(\alpha) \geq 0$.

This ring has the maximal ideal

$$M = \{\alpha : v(\alpha) > 0\}.$$

Indeed, if we take any other element in the ring, $v(\alpha) = 0$. Then $v(\alpha^{-1}) = 0$ as well, so α is a unit of R. In other words, all elements not in M are units, so M must be a maximal ideal.

M is also a principal ideal of the form xR, where x is any element where $v(x) = 1$. Indeed, if $y \in R$ and $v(y) = n$, then $v(x^{-n}y) = 0$. Thus $u = x^{-n}y$ is a unit, so $y = x^n u$, meaning that $y \in x^n R = M^n$. Thus, the ideals of R are either the zero ideal or of the form R^n, and by the same logic, any element y of a proper R-submodule of K is of the form $x^k R$, and thus $N = x^k R$ (for the minimal $k \in \mathbb{Z}$) or K (if no such k exists).

This also tells us that there is no ring R' such that $R < R' < K$. If there were such an R', it would be an R-submodule of K, and it must be of the form $x^k R$ for some $k \geq 0$ (or else x^k squared wouldn't be in the ring). But R is contained in R', so $k = 0$.

Definition 112

A **local ring** R is a Noetherian ring with exactly one maximal ideal M, where M is not the zero ideal.

In particular, a valuation ring is a local ring.

Proposition 113

Let R be a local domain with maximal ideal M. Then M is a principal ideal if and only if R is a valuation ring.

Proof. We've already shown the backwards direction above. For the forward direction, first assume that the maximal ideal in our local ring $M = xR$ is a principal ideal. Let y be a nonzero element of R: by the Nakayama Lemma, we know that the powers of x that divide y are bounded, so we can write $y = x^k u$, where x does not divide u. But u is not divisible by x, so it must be a unit in the ring: we can then define $v(y) = k$. So we've now defined a valuation, and it just remains to check the other conditions in the definition (which are not too hard). □

With this, we can move on to the concept of a local ring at a point:

Definition 114

Let X be an affine variety, and let p be a point of X. Then the **local ring of X at p** is

$$A_p = \{\alpha \in K : \alpha \text{ regular at } p\}.$$

In other words, there exists an affine neighborhood $\operatorname{Spec} A$ of the point p in X, such that $\alpha \in A$.

To explain why this is indeed a local ring, we need to show that there is only one maximal ideal. The maximal ideal at a point p of X is

$$\mathfrak{m}_p = \{\alpha : \alpha \text{ regular at } p, \alpha(p) = 0\}.$$

It suffices to show that if $s \in A_p$ and $\alpha(s) \neq 0$, then α is a unit, so $\alpha^{-1} \in A_p$ as well. We know that $\alpha^{-1} \in A[\alpha^{-1}]$, and $\operatorname{Spec} A[\alpha^{-1}]$ is $\operatorname{Spec} A$ with the zeros of α removed. Thus, $p \in \operatorname{Spec} A[x^{-1}]$, so α^{-1} is regular at p, as desired.

We'll now move on to our next topic – **normalization**.

Definition 115

Let A be a Noetherian domain with a fraction field K of characteristic zero. Then the **normalization** of A, denoted $A^{\#}$, is the set of $\alpha \in K$ that are integral over A. A is **normal** if it's equal to its normalization $A^{\#}$.

It is clear that $A \subset A^{\#} \subset K$, and we have some more structure here:

Theorem 116

$A^{\#}$ is a finite A-module. In other words, the map $X^{\#} \to X$ is an integral morphism.

Proof sketch. The proof is involved, and one important object that will come up later in the class is the **trace**.

Definition 117

Let L be a finite field extension of K, and let $\beta \in K$. Multiplication by β is a linear operator on L – the **trace** of β is the trace of this linear operator.

We can use this trace to define a bilinear form on L, sending α, β to $\text{tr}(\alpha\beta)$. This is an element of A (this is easy to verify because the trace of an operator is basis-independent). From here, the main ideas of the proof are to assume A is normal, use a basis for L such that all elements are integral, and then consider the map of L into K^n via the operator

$$T(\beta) = (\langle v_1, \beta \rangle, \cdots, \langle v_n, \beta \rangle).$$

This is injective, and it maps the integral closure B of A (in L) to A^n, which means that B is isomorphic to a submodule of A^n, which is finitely generated because A is Noetherian. □

Let's show an example of this (the most important case is that of curves):

Example 118

Let $A = \mathbb{C}[x, y]/(y^2 - x^3)$, and let $z = \frac{y}{x}$.

z is integral over A, since $z^2 = \frac{y^2}{x^2} = \frac{x^3}{x^2} = x$ (where we're using x and y for the residues in A as well). In fact, we also have $y = z^3$, so the normalization $A^\#$ is just polynomials $\mathbb{C}[z]$. In other words, we've "straightened out the cusp" to create the affine line! We do need the following result to show that we've found our normalization, though:

Lemma 119

Any unique factorization domain is normal – in particular, polynomial rings over a field, such as $\mathbb{C}[z]$, are normal.

We can show this by showing that any element $\beta = \frac{r}{s}$ is integral and writing out its defining equation.

Example 120

Consider the ring $A = \mathbb{Z}[\delta]$, where $\delta^2 = -3$.

As we might know from 18.702, this isn't an integrally closed ring: we get the nice properties by using $A^\# = \mathbb{Z}[\eta]$, where $\eta = \frac{1+\delta}{2}$.

Definition 121

A local ring A has **dimension** 1 if the only prime ideals of A are 0 and the maximal ideal M.

For example, the valuation v_p at a point p is a one-dimensional local ring.

Theorem 122

Let A be a local ring. Then A is one-dimensional and normal if and only if it is a valuation ring.

Proof. The backwards direction is easy, and the forward direction is trickier. Our plan will be to show that the maximal ideal M of A is a principal ideal, and then we can use Proposition 113. Let $x \in M$: the radical ideal $\text{rad}(xR)$ is the intersection of some prime ideals, but there are only two prime ideals in our local ring, and one of them is the zero ideal. Thus $\text{rad}(xR) = M$, which means that $M^r \subset xR$ for some $r > 0$.

Let r be the minimal such possible exponent: then M^{r-1} is not contained in xR, but M^r is, and we can pick an element of y that is in M^{r-1} but not xR. So now $yM \subset M^r \subset xR$, and now we can consider the fraction $z = \frac{y}{x}$. This is not in R (because y is not in xR) but zM is in R (because yM is in xR).

But now notice that zM is an ideal of R, because it's closed under multiplication by R and addition. Since there is only one maximal ideal, either zM is contained in M, or zM is the whole ring. Let's write out those two cases:

(1) If zM is the whole ring R, then we can let $t = \frac{1}{z}$, and we know that $M = tR$ (by multiplying through by t). Thus t is actually in R (because M contains $t \cdot 1$), and now we've shown that M is a principal ideal, so we're done.

(2) if $zM \subset M$, we can apply Nakayama: say that $v_1, \cdots, v_k$ generate the ideal M. Then we can write

$$zv_i = \sum_j p_{ij} v_j,$$

where the p_{ij} are in R. Therefore, $(v_1, \cdots, v_k)$ is an eigenvector of P with eigenvalue z, which means that $p(z) = 0$. But this means that z is integral, which is a contradiction (since z is not in R). □

Fact 123

We will also include the proof of Krull's theorem here, now that we have enough tools. Recall that we want to show that any closed subset $V_X(\alpha)$, corresponding to the zero set of α in an affine vareity X of dimension d, has irreducible components of dimension $(d-1)$.

Proof of Krull's Theorem. We know that the dimension of any component C is less than d, because C is a proper closed subset of X. Assume for the sake of contradiction that $\dim C < d - 1$.

Let A be the coordinate algebra of X, and let B be its normalization. Then $Y = \text{Spec}\, B$ has dimension d (because A and B have the same fraction field, thus the same transcendence degree), and the morphism $Y \to X$ is an integral morphism by definition. This map is also surjective, because every point in X is a maximal ideal $\mathfrak{m}$ in A, and the extension $\mathfrak{m}B$ in B is not the unit ideal (this is a lemma from earlier in the class), so it corresponds to a point in Y. Since we have an integral morphism, closed sets are mapped to closed sets, and there is some irreducible component of $V_Y(\alpha)$ which maps to our irreducible component C of $V_X(\alpha)$. Now the morphism $D \to C$ is an integral morphism as well (it's defined the same way as the one from Y to X), so D and C have the same dimension, meaning that it suffices to prove Krull's theorem for the case where A is **normal** (because we can replace X with Y and have the same dimensionality).

Let the zero locus of α be written as $C \cup V$. Similarly, as our earlier attempt, we want V to be empty, so we'll localize by finding an element s that is identically zero on V but not on C. Then V_s is empty (because s is zero everywhere), but C_s is not empty. We've assumed here that A (and therefore X) is normal, so the localization X_s is normal as well. Since X_s and C_s have the same dimensions as X and C, respectively, we can use the localizations instead, and this means that we can assume α's zero locus is **irreducible**.

Now we can say that C is both the zero locus of P and of α. This means that $P = \text{rad}\, P = \text{rad}(\alpha A)$, which means that $P^n \subseteq \alpha A$ for sufficiently large n. We're assuming that C has codimension at least 2 (in other words, dimension at most $d-2$). By definition, this means that $C \subset Z$ for some subvariety Z of X with codimension 1 – say that this corresponds to the prime ideal Q. Now $P \supset Q$, but α is not in Q (since α does not vanish on all of X), so αA is not contained in Q.

Lemma 124

With the notation above, there is an element $\gamma \in A$ such that γ is not in αA but $P\gamma$ is contained in αA.

Proof of lemma. Pick some $\beta \in Q$. We have a Noetherian domain here, so by corollaries of the Nakayama Lemma, we know that there is some k such that α^k is the largest power of α dividing β. Now $\gamma = \frac{\beta}{\alpha^k}$ vanishes on $Z \setminus C$, which is a dense subset of Z, which means that γ vanishes on all of Z. So now we have a preliminary candidate: $\gamma \in Q$ and α does not divide γ.

Now to show that we can find a γ such that $P\gamma$ is contained in αA, we know from above that $P^n \subseteq \alpha A$ for some n and thus $P^n Q \subseteq \alpha A$ – let r be the smallest integer such that $P^r Q \not\subseteq \alpha A$. (This exists because we found a preliminary γ above) Then there is some element of this ring $P^r Q$ that is not in αA, and that γ satisfies the conditions that we want. □

Taking this γ, we can consider the element $\delta = \frac{\gamma}{\alpha}$. We know that $P\delta \subset P$, because α is in P, but $\gamma \notin \alpha A \implies \delta \notin A$ by construction. $P\delta$ then vanishes on $Z \setminus C$, which is a dense subset of Z, which means that $P\delta$ vanishes on all of Z. Therefore, $P\delta \subset P$, meaning δ is integral over A. By normality, this means that $\delta \in A$, which is a contradiction. □

20 April 3, 2020

Definition 125

A **curve** X is a variety of dimension 1.

The proper closed subsets of a curve are the finite subsets (by Krull's theorem).

Definition 126

Writing a curve X as $\operatorname{Spec} A$, if A is normal, then X is a **smooth** affine curve. In general, X is **smooth** if it can be covered by smooth affine curves.

In general, if X is not smooth, we can take the normalization of A, and then $X^\#$ will be a smooth curve. (We need to normalize each of the affine open sets of a covering, and we need to check that this is consistent.) Then the morphism $X^\# \to X$ is integral (because every element of $A^\#$ is integral over A), so we have a **surjective** morphism.

Example 127

A good way to think about this is that $X^\#$ removes singular points from X: for example, a cusp will be straightened out.

Proposition 128

Local rings of smooth curves are valuation rings.

By definition, smooth curves are dimension 1. Remember that the local ring at each point is the set of rational functions that are regular, so we just need to know that (assuming that X is affine here) the local ring is also normal. But we won't focus too much on the details here! We should just think of the local ring at each point as corresponding to the valuation v_p which evaluates the multiplicity of the zero or pole.

We'll move along to talking about **constructible sets**: throughout this discussion, X is a variety.

Definition 129

A **locally closed set** L is an intersection of the form $Y \cap U$, where Y is closed and U is open.

Example 130

Let $X = \mathbb{A}^2$, and let L be a line with one point p deleted. Then this is the intersection of L with the complement of p, so this is a locally closed set.

Clearly, open and closed sets are locally closed.

Proposition 131

L is locally closed if and only if it is a closed subset of an open set, which is also equivalent to it being an open subset of a closed set.

Proof. We can replace $Y \cap U$ with $Y \cap (U \cap Y)$ to get an open subset of a closed set, and we can replace it with $U \cap (U \cap Y)$ to get a closed subset of an open set. In general, we can find the minimal open and minimal closed sets containing L (which exist by looking at chains of closed sets). □

Definition 132

A **constructible set** S is a finite union of locally closed sets.

Example 133

Take $X = \mathbb{A}^2$, and let S be the union of the complement of the y-axis with the origin. Then S is constructible.

Lemma 134

Suppose that S is constructible. Then we can write $S = L_1 \cup \cdots \cup L_k$, such that $L_i = Y_i \cap U_i$ and the Y_i are **irreducible and distinct**.

Proof. We decompose each closed component Y_i into its irreducible components, and then we "merge" parts that share the same closed set – this preserves the property that all L_i are locally closed. □

Theorem 135

Let $\mathcal{S}$ be the smallest family of subsets such that

- open sets are in $\mathcal{S}$,
- finite unions and intersections are in $\mathcal{S}$ ($S_1 \cup S_2$ and $S_1 \cap S_2$ are in $\mathcal{S}$ for $S_1, S_2 \in \mathcal{S}$).
- complements $S^C = X - S$ are in $\mathcal{S}$ for any $S \in \mathcal{S}$.

Then $\mathcal{S}$ is the family of constructible sets.

All constructible sets can be formed from open sets. To prove the remainder of the result, note that the union and intersection of two constructible sets is constructible – if $S_1 = L_1 \cup \cdots \cup L_k$ and $S_2 = M_1 \cup \cdots \cup M_n$, we can write $S_1 \cap S_2 = \bigcup_{i,j} L_i \cap M_j$. This is because the intersection of locally closed sets is locally closed. We can similarly show that complements are also constructible, so we can show closure.

Theorem 136

Suppose $f : Y \to X$ is a morphism, and S is a constructible subset of Y. Then $f(S)$ is a constructible subset of X.

Conversely, the inverse image $f^{-1}(T)$ of a constructible set $T \in X$ is constructible (but this is easier; it follows from continuity of morphisms).

Proof. (This basically consists of reducing the problem until there is nothing left to prove.) Since we're working with Noetherian rings, we can prove this by **Noetherian induction** – by the descending chain condition, we just need to show that if a statement is true for all closed proper subsets of Y, it is true for Y as well (and the same is true for X, which is also closed). In other words, we can assume this statement is already true if S is contained in a proper closed subvariety of Y or $f(S)$ is contained in a proper closed subvariety of X.

We can now proceed with the proof. If Y_1 is a proper closed subvariety of Y, we can write the constructible set S as $(S \cap Y_1) \cup (S \cap (Y \setminus Y_1)))$, and we can check that both of these pieces are constructible subsets. But then we've shown the condition for $S \cap Y_1$ by Noetherian induction already, so we can replace Y with $Y \setminus Y_1$ finitely many times, which lets us assume that Y is a nonempty **open subvariety**. By basically the same logic, we can replace X with a nonempty open subvariety.

It suffices to show that a locally closed set satisfies this condition (because S is a finite union of such sets), and we can write this set as $C \cap U$, where C is closed and irreducible and U is open. Now $Y = C \cap (Y \setminus C)$, but $Y \setminus C$ does not intersect C, so we can assume that Y is just C. In other words, $S = C \cap U = Y \cap U$, and now we can assume Y is just U – we've now reduced the problem to the case where $S = Y$.

Remember that we showed above that we can replace X and Y with open subvarieties, so we can write $Y = \operatorname{Spec} B$ and $X = \operatorname{Spec} A$. We now have an associated homomorphism $\phi : A \to B$ – this homomorphism is **injective** because the kernel corresponds to a proper closed subset (which we've already removed from the picture in an earlier simplification).

To finish, Noether Normalization tells us that we can pick $s \in A$ such that B_s is a finite module over $A_s[y_1, \cdots, y_k]$. We have surjective maps from Y_s to $\operatorname{Spec} A_s$, as well as from $\operatorname{Spec} A_s$ to X_s (because we have a finite module), which means that there is a surjective map from Y_s to X_s. Once again replacing X and Y with the open subsets X_s and Y_s gives us a surjective map, which indeed means that our image is constructible, as desired.

□

We'll move on again to a nice topic: **analyzing with smooth curves**. Suppose that C is a smooth affine curve, q is a point of C, and $C' = C - \{q\}$ is its complement. We should think of q as a "limit point."

Proposition 137

C is the closure of C' in the Zariski and in the classical topology.

Proof. This is true in the Zariski topology because C is irreducible (so we can't break it up further, meaning C' is not closed). In the classical topology, C''s closure must be either C' or C because C is closed – if it were C', then q would be open, but it is also closed. Thus, q would be an isolated point, and no such points exist. □

Theorem 138

Let X be a variety, and let S be a constructible set in X. Then S is closed if and only if for all morphisms $f : C \to X$, we have that $f(C') \subset S \implies f(C) \subset S$.

In other words, we can test if S is constructible by testing curves. Think of q as a limit point of C': this says that if S contains all of its limit points, it's closed.

Proof sketch. The idea is that we need to "find enough curves:" we need to be able to test any point p in the closure $\overline{S}$ to see if it is in S. To do this, we map in a curve C intersecting S at this point p, and use that as the removed point. Here's the main lemma that we need:

Lemma 139

Let S be a constructible subset of X, and let $p \in \overline{S}$ be a point. Then there exists a curve C and a point q such that there exists a morphism $f : C \to X$ with $f(C') \subset S$ and $f(q) = p$.

Proof of lemma. Let $S = L_1 \cup \cdots \cup L_k$, where $L_i = Y_i \cap U_i$; we can assume that the Y_i are irreducible by Lemma 134. Each L_i is then some dense open subset of Y_i, meaning the closure of each L_i is Y_i. Thus, the closure of $\overline{S}$ is the union of Y_1 through Y_k, meaning any point p is in Y_i for some i.

Therefore, there exists an **irreducible** closed set Y that contains p, as well as a nonempty open subset $V \in Y$ such that $V \subset S$ (this is basically the corresponding locally closed set L_i). We can assume that Y is affine with $Y = \operatorname{Spec} B$ (by taking X to be affine – this is okay because we have a Noetherian space). If $\dim Y = 0$, then Y is just a single point, so $p \in S$. Otherwise, we can reduce to $\dim Y = 1$ by Krull's theorem: this is because we can take a zero locus from an element $\beta \in B$ such that $\beta(p) = 0$, but β is not identically zero on any component of Z except p itself, and consider $\beta \cap Y$. By construction, this zero locus is closed and nonempty, and because it contains p, at least one of these components of the zero locus contains p, and we can replace Y by one of those components – this will decrease the dimension by at least 1.

So we can now say that $\dim Y = 1$. Let $C = Y^{\#}$ be a smooth curve: we know that $\pi : Y^{\#} \to Y$ is surjective (because we have an integral morphism), so there exists a $q \in C$ such that $\pi(q) = p$. We delete the (finite set of) other points in the fibre over p by picking an affine open subset, and this shows that we indeed have the morphism that we desire. □

Using this curve to "test" closure gives us the result – if such a limit point p were not in S (meaning that S is not closed), then we would not have that $f(C') \subset S \implies f(C) \subset S$. □

Corollary 140

If S is a **constructible** subset of X, then it is closed in the Zariski topology if and only if it is closed in the classical topology.

The classical topology is finer, so anything closed in the Zariski topology is closed in the classical topology. The converse comes from the "testing curve" idea above.

We'll finish today with two more results about **projective space being proper**:

Theorem 141

Let Z be a variety, and let Y be a closed subset of $Z \times \mathbb{P}^n$. Let $\pi : Z \times \mathbb{P}^n \to Z$ be the projection map. Then $\pi(Y)$ is closed in Z.

Proof. This is true in the classical topology because $\mathbb{P}^n$ is compact, so projection takes closed sets to closed sets. More explicitly, take a sequence of points p_i in $\pi(Y)$ converging to a limit point – for each one, we can pick a point

(z_i, x_i) in $Y \subset Z \times \mathbb{P}^n$ that maps to it under π, and then we can pick a subsequence of $\{x_i\}$ with a limit point x in $\mathbb{P}^n$, because of compactness. The z_i also converge to some z (the same limit point that the p_i converge to). Now (z, x) is in Y (because Y is closed and therefore contains all of its limit points), and therefore its image p, which is the limit point of the p_is, is also in $\pi(Y)$.

But now images of constructible subsets are constructible by Theorem 136, and Y is constructible (because it is closed). Because $\pi(Y)$ is closed in the classical topology, it is closed in the Zariski topology. □

Corollary 142

Let $f : X \to Z$ be a morphism of **projective varieties**, and Y be closed in X. Then $f(Y)$ is closed in Z.

Proof. Use the graph Γ_f, which consists of the points $(x, f(x)) \subset X \times Z$. This is isomorphic to X, and so Y is isomorphic to a closed subset of $\Gamma_f \subset X \times Z$. Now getting the values for $f(Y)$ is a projection map, and $X \subset \mathbb{P}^n \implies \Gamma_f \subset \mathbb{P}^n \times Z$, meaning that we can use the above result. □

21 April 6, 2020

We'll be spending the next week or so on **modules** – the first topic here is the structure sheaf of a variety. Let X be a variety (everything here is defined with respect to X) – we can define the **category (opens)**, where the objects of the category are open subsets of X and the morphisms are inclusions. (There is a morphism from $V \to U$ if $V \subset U$, and no morphism if $V \not\subset U$.)

Definition 143

The **structure sheaf** $\mathcal{O}$ is the functor from (opens) to (algebras). Specifically, if U is open, then $\mathcal{O}(U)$ is the algebra of rational functions α that are regular on U.

Recall that if α is in a function field K, then α is **regular** at a point $p \in X$ if there exists an affine open set such that $U = \operatorname{Spec} A$, $p \in U$, and $\alpha \in A$. Now if we have an inclusion $V \to U$ of open sets, then α being regular on U implies that it is regular on V – this means that $\boxed{V \subset U \implies \mathcal{O}(U) \subset \mathcal{O}(V)}$, and that's what makes the structure sheaf into a **functor**. (Notice that the arrows are reversed here, so this is **contravariant**.)

Suppose we want to check that a rational function is regular on some open set or on X – we don't want to check every affine open set.

Lemma 144

Let $U = \operatorname{Spec} A$ be an affine open set in X that is covered by other affine sets $\{U^i = \operatorname{Spec} A_i\}$ (we can assume this set is finite). Then $A = \bigcap A_i$.

In other words, the rational functions on $\operatorname{Spec} A$ are just the elements of A, and checking whether functions are regular only requires a specific open covering, not all of them.

Proof. Note that **being covered** here means that the U^i are all subsets of U, so having maps $U^i \to U$ means we have maps $A \to A^i$ – this means that every element of A is in $\bigcap A^i$.

We know that the localizations of U form a basis for the topology. Thus, we can always assume our open coverings are localizations, and write $U^i = \operatorname{Spec} A[s_i^{-1}]$, where $s_i \in A$ are nonzero for all $1 \leq i \leq k$. The U^is cover U, so there is no point where the s_is are all zero – that means that $s_1, \cdots, s_k$ generate the unit ideal A. Let α be an element of the intersection $\bigcap A_i$: since $\alpha \in A[s_i^{-1}]$, we know that $s_i^n \alpha \in A$ for some n. There are finitely many s_is, so we can use the same n for everything. Since $\{s_i\}$ generate the unit ideal, so do the powers $\{s_i^n\}$, and thus we can write

$$1 = \sum_i a_i s_i^n \implies \alpha = \sum a_i s_i^n \alpha \in A,$$

and thus every element of $\bigcap A_i$ is in A. □

Lemma 145

Let U be any open set in X, and let $\{U^i = \operatorname{Spec} A_i\}$ be an affine open cover of U. Then $\mathcal{O}(U) = \bigcap \mathcal{O}(U^i)$.

The lemma tells us that the affine open sets are the important ones here. This is basically the same as the argument we made above – we'll omit the proof.

Proposition 146

$\mathcal{O}$ has the **sheaf property**: $\mathcal{O}(U) = \bigcap \mathcal{O}(U^i)$ if $\{U^i\}$ is an affine cover of U.

The structure sheaf is a functor, and we want it because the variety X is not just a topological space – all one-dimensional curves are homeomorphic, but we do care about what the functions are.

We'll now move on to the definition of $\mathcal{O}$-modules: let **(affines)** be the subcategory of (opens), where the objects are the affine open subsets of X, and the morphisms are localizations. This second part is not so obvious: what it says is that a morphism $V \to U$ in (opens) – that is, an inclusion – will be a morphism in affines if (1) $U = \operatorname{Spec} A$ is affine and (2) V is a localization of U, meaning that $V = U_s = \operatorname{Spec} A[s^{-1}]$ for some $s \neq 0$ in A.

The reason we like to restrict to (affines) is that we know what localizations look like, and we don't have very much to say about inclusions of open sets that aren't localizations.

Definition 147

An **$\mathcal{O}$-module** is a functor $\mathcal{M}$ from (affines) to (modules).

So in an $\mathcal{O}$-module, we assign a module to each affine open set of X, which should be compatible with the localization. Specifically, if $U = \operatorname{Spec} A$ is an affine open set, the $\mathcal{M}(U)$ is an $\mathcal{O}(U)$-module, where $\mathcal{O}(U) = A$. Since our maps look like localizations, we can explain this more explicitly: if U_s is a localization of U, then $\mathcal{M}(U_s)$ is the localization of $\mathcal{M}(U)$, which we denote as $\mathcal{M}(U)_s$. In particular, if $\mathcal{M}(U)$ is an A-module M, then $\mathcal{M}(U)_s$ is just the localization M_s.

Definition 148

Elements of $\mathcal{M}(U)$ are **sections** of $\mathcal{M}$ on U.If $V \to U$ is a map in (affines), so $V = U_s$, then the image of a section $m \in \mathcal{M}(U)$ is the **restriction** of m to V.

By custom, we denote the restriction to V by the same letter m.

Example 149

In the free $\mathcal{O}$-module $\mathcal{O}^n$, the sections are vectors of sections of $\mathcal{O}$.

Example 150

If $X = \operatorname{Spec} A$ is affine, then $\mathcal{O}$-modules are A-modules and vice versa.

Example 151

Let $p \in X$ be a point. The **residue field module** κ_p is defined as follows: if U is an element of (affines) and $U = \operatorname{Spec} A$, such that $p \in U$, then A has a residue field $\kappa(p)$, and we define $\kappa_p(u) = \kappa(p)$. If $p \notin U$, then we define $\kappa_p(U) = 0$.

Definition 152

An **ideal** I of $\mathcal{O}$ is an $\mathcal{O}$-submodule of $\mathcal{O}$. A **maximal ideal** at a point $p \in X$ is $\mathfrak{m}_p(U)$ if $U = \operatorname{Spec} A$ and $p \in U$ and $\mathcal{O}(U)$ if $p \notin U$.

Definition 153

Let $\mathcal{M}$ and $\mathcal{N}$ be $\mathcal{O}$-modules. A **homomorphism** $\phi : \mathcal{M} \to \mathcal{N}$ is defined as follows: for every $U \in \operatorname{Spec} A$, we know that $\mathcal{M}(U)$ and $\mathcal{N}(U)$ are $\mathcal{O}(U)$-modules, and we have the homomorphism $\phi : \mathcal{M}(U) \to \mathcal{N}(U)$ of $\mathcal{O}(U)$-modules, such that whenever U_s is the localization of U, $\phi(U_s)$ is the localization of $\phi(U)$.

It seems annoying to define homomorphisms on every affine open set, but all of the baggage takes care of itself:

Example 154

Let $\phi : M \to N$ be a homomorphism. Then we have a kernel K, a cokernel C, and an image – we know that the kernel K should be an $\mathcal{O}$-module, which means we need to understand $K(U)$ for any affine U. But we can just say that $K(U)$ is the kernel of the homomorphism $\phi(U) : \mathcal{M}(U) \to \mathcal{N}(U)$.

We just need to make sure that localization is compatible:

Lemma 155

If $\phi : M \to N$ is a homomorphism of A-modules for a domain A, and $s \in A$ is a nonzero element, then we have a homomorphism $\phi_s : M_s \to N_s$ by localization. Then $\ker \phi_s = (\ker \phi)_s$.

Definition 156

An sequence of $\mathcal{O}$-modules $\cdots M \xrightarrow{f} N \xrightarrow{g} P \cdots$ is **exact** if for any two adjacent maps f, g as above, we have $\operatorname{Im} f = \ker g$.

Example 157

We have an exact sequence

$$0 \to \mathfrak{m}_p \to \mathcal{O} \to \kappa_p \to 0,$$

which means that the kernel of the homomorphism $\mathcal{O} \to \kappa_p$ is $\mathfrak{m}_p$ and that we have an injective map $\mathfrak{m}_p \to \mathcal{O}$.

22 April 8, 2020

We'll continue to discuss modules today. Recall that $\mathcal{O}$ is the **structure sheaf** on a variety X, where $\mathcal{O}(U)$ is a regular function on U. Then an **$\mathcal{O}$-module** $\mathcal{M}$ is a (contravariant) functor associates a module to every affine open set: if $U = \operatorname{Spec} A$ is an affine open set, then $\mathcal{M}(U)$ will be an $\mathcal{O}(U)$-module, which is just an A-module.

We defined the category (affines) last time; the morphisms in this category are the **localizations** $U_s \to U$ for some nonzero $s \in A$. In order for this to be valid, we just need $\mathcal{M}(U_s)$ to be the localization of $\mathcal{M}(U)$. And again, if we have an affine open set $\operatorname{Spec} A$, then $\mathcal{M}(U_s) = M_s$ is an A_s-module.

Our goal is to extend $\mathcal{M}$ to all opens (instead of just affines) – our question is to ask for **global sections** of $\mathcal{M}$, which are the elements of $\mathcal{M}(X)$. They should be given by sections on an open covering $\{U^i = \operatorname{Spec} A_i\}$ which agree on the overlap, so any element $\mathcal{M}(U)$ should be defined with a vector $(m_1, \cdots, m_k)$, where $m_i \in \mathcal{M}(U^i)$, such that the restrictions of m_i, m_j to the smaller open set U^{ij} are equal.

But we know that $U^{ij} = U^i \cap U^j$ is the intersection of affine open sets, so it is affine open as well: it's $\operatorname{Spec}[A_i, A_j]$. This **does not always** give us a map from $\mathcal{M}(U^i)$ to $\mathcal{M}(U^{ij})$ unless U^{ij} is actually a localization, but it's not such an important point.

In order for this to work out, we want the following sequence to be exact:

$$\boxed{0 \to \mathcal{M}(U) \to \prod_i \mathcal{M}(U^i) \to \prod \mathcal{M}(U^{ij})},$$

where the first arrow tells us that the map in the second arrow from $\mathcal{M}(U) \to \prod_i \mathcal{M}(U^i)$ is injective (and describes the vector $(m_1, \cdots, m_k)$). Meanwhile, we should think of the last space $\prod \mathcal{M}(U^{ij})$ as the space of $k \times k$ matrices: then we want our map β to send $(m_1, \cdots, m_k) \to (z_{ij})$, where

$$z_{ij} = m_j - m_i$$

when restricted to U^{ij}. When these conditions are satisfied, this yields the **sheaf property**.

Theorem 158

Let $\mathcal{M}$ be an $\mathcal{O}$-module. Then $\mathcal{M}$ extends uniquely to a functor $\mathcal{M}$ from (opens) to (modules) – such that $\mathcal{M}$ has the sheaf property. In addition, any homomorphism of $\mathcal{O}$-modules $f : \mathcal{M} \to \mathcal{N}$ extends uniquely.

This proof is long – we need to check that our map β above makes sense, which amounts to showing that when we ave two affines with one contained in the other, there exists a map. We should read about this on our own.

Note that in our exact sequence, we can just take the product $\prod_{i<j} \mathcal{M}(U^{ij})$, because U^{ij} and U^{ji} are the same and $U^{ii} = U^i$.

Example 159

Suppose $X = U^0 \cup U^1$ is covered by those two affine open subsets. Then the sheaf axiom is equivalent to the following sequence being exact:

$$0 \to \mathcal{M}(U) \to \mathcal{M}(\mathcal{U}') \times \mathcal{M}(U^1) \to \mathcal{M}(U^{01}).$$

If we say that $X = \mathbb{P}^1$ is the projective line, $U^0 = U^1$ are the standard affine open sets in X, and $U^0 = \operatorname{Spec} A_0$ and $U_1 = \operatorname{Spec} A_1$. Then $A_0 = \mathbb{C}[u]$, where $u = \frac{x_1}{x_0}$, and $A_1 = \mathbb{C}[v] = \frac{x_0}{x_1} = u^{-1}$. Thus

$$U^{01} = \operatorname{Spec} \mathbb{C}[u, u^{-1}]$$

which is the Laurent polynomials in u. So if we have an $\mathcal{O}$-module, we can get an A_0-module $\mathcal{M}_0$ and an A_1-module $\mathcal{M}_1$. Now the $\mathcal{M}_i$ are the sections of $\mathcal{M}(U^i)$, and we must have

$$\mathcal{M}(U^{01}) = \mathcal{M}_0[u^{-1}] = \mathcal{M}_1[v^{-1}].$$

On the other hand, if we have an A_0-module $\mathcal{M}_0$ and an A_1-module $\mathcal{M}_1$, and suppose there's an isomorphism $\theta : \mathcal{M}_0[u^{-1}] \to \mathcal{M}_1[v^{-1}]$ An $\mathcal{O}$-module, we should have compatibility from localization, and thus we get an $\mathcal{O}$-module $\mathcal{M}$.

Example 160

Suppose $\mathcal{M}_0$ is a free A_0-module with basis B_0, and $\mathcal{M}_1$ is a free B_1-module with basis B_1.

When we do our localizations, $M_0[u^{-1}]$ and $M_1[v^{-1}]$ are free A_{01}-modules with bases B_0, B_1, and they are isomorphic. So θ is some invertible matrix P with entries in A_{01} – note that we can chagne B_0 by an invertible A_0-matrix Q_0 and and B_1 by an invertible A_1-matrix to Q_1, so here P determines $\mathcal{M}$.

Theorem 161 (Birkhoff-Grothendieck)

Given an invertible A_{01}-matrix P (this is the Laurent polynomial ring), there exist an invertible A_0-matrix Q_0 and an invertible A_1-matrix Q_1 such that $Q_0^{-1}PQ_1$ is diagonal.

We know that the units of A_0 are cu^k for some $k \in \mathbb{Z}$, This means that $Q_0^{-1}PQ_1$ looks like a diagonal matrix with entries of the form $u_i^{e_i}$, where all e_i are integers. We'll prove this later on with cohomology.

Corollary 162

Suppose that $X = \operatorname{Spec} A$ is affine. Then an $\mathcal{O}$-module is equivalent to an A-module.

Proof. If we're given an $\mathcal{O}$-module $\mathcal{M}$, then $M = \mathcal{M}(X)$ is an A-module. The localizations form a basis, so it's enough to just define $\mathcal{M}(X_s) = M_s$. □

So the localizations are enough in this case where x is just an (affine)!

Proposition 163 (Coherence property)

If Y is any open set, and $s \in \mathcal{O}_y$ is a regular function, then Y_s is $Y \setminus \{(\text{zeros of } s\}$. Then $\mathcal{M}(Y_s)$ is the localization of Y.

Remember that the definition of a module $\mathcal{M}$ means it has to be given on affines – thus we know that this is true by definition for an affine open set Y. This property extends that fact to opens in general! We then call an $\mathcal{O}$-module M that has been extended to all opens a **quasicoherent sheaf**. Note that working with the extended module is harder – we like working with (affines) and localizations, because many operations on modules are compatible with localizations, while very few operations work on all open sets. The only notable exception to this is the kernel – for example, the cokernel doesn't work:

Example 164

Suppose $f : \mathcal{M} \to \mathcal{N}$ is a homomorphism of $\mathcal{O}$-modules. Then $c = \operatorname{coker} f$ is the module such that $C(U)$ is the cokernel of the map $\mathcal{M}(U) \to \mathcal{N}(U)$ for an affine open set $U = \operatorname{Spec} A$, but not otherwise.

23 April 10, 2020

Recall that for an $\mathcal{O}$-module $\mathcal{M}$ for a variety X extends uniquely to a functor on all opens with the sheaf property.

Lemma 165

The sections of $\mathcal{M}$ on the empty set are $\{0\}$.

This is more semantic than anything else – we're bringing it up because we've been avoiding the empty set.

Example 166

A **tensor product** is an example of an $\mathcal{O}$-module.

Recall that that if we have two A-modules M, N, we can construct an A-module $M \otimes N$ generated by tensors of the form $m \otimes_A n$ (with $m \in M, n \in N$), with the bilinear relations

$$(m_1 + m_2) \otimes n = m_1 \otimes n + m_2 \otimes n, \quad m \otimes (n_1 + n_2) = m \otimes n_1 + m \otimes n_2, \quad am \otimes n = m \otimes an.$$

One special case, where M and N are free modules with bases $(m_1, \cdots, m_r)$ and $(n_1, \cdots, n_s)$, yields a tensor product $M \otimes N$ which is a free module with basis $\{m_i \otimes n_j\}$. We can think of $M = A^r$ as forming the rows and $N = A^s$ as forming the columns of an $r \times s$ matrix.

Definition 167

Let $\mathcal{M}$ and $\mathcal{N}$ be $\mathcal{O}$-modules. Then the **tensor product** is defined by its sections on U:

$$\mathcal{M} \otimes_{\mathcal{O}} N(U) = M(U) \otimes_{\mathcal{O}(U)} N(U).$$

(In other words, the scalars of the tensor product are the regular functions on U.) This is a valid definition, because taking any nonzero element $s \in A$ yields

$$(M \otimes_A N)_s \cong M_s \otimes_{A_s} N_s.$$

Definition 168

Let $f : Y \to X$ be a morphism, and let $\mathcal{N}$ be an $\mathcal{O}_Y$-module. The **direct image** $f_*\mathcal{N}$ is an $\mathcal{O}_X$-module defined as follows: suppose $U = \operatorname{Spec} A$ is an affine open set in X, meaning that $f^{-1}U$ is some open set in Y (not necessarly affine). Then the sections of $f_*\mathcal{N}$ on U are the sections of $\mathcal{N}$ on $f^{-1}U$.

We can write this with less words as

$$(f_*\mathcal{N})(U) = \mathcal{N}(f^{-1}U).$$

Lemma 169

The direct image $f_*\mathcal{N}$ is an $\mathcal{O}_X$-module.

We prove this using the **coherence property**.

Example 170

Suppose $Y \xrightarrow{i} X$ is an inclusion of a **closed** subvariety Y into X. Then $i^{-1}(U) = U \cap Y$.

Then we should have

$$f_* \mathcal{N}(U) = \mathcal{N}(U \cap Y).$$

We're mentioning this special case because it's called the **extension of $\mathcal{N}$ by zero**. This is because any open set $U \in X$ that doesn't intersect Y yields

$$(f_* \mathcal{N})(U) = \mathcal{N}(\varnothing) = 0.$$

Example 171

Suppose $Y \xrightarrow{j} X$ is an inclusion of an **open** set.

Inverse images are intersections again, so we still have

$$(f_* \mathcal{N})(U) = \mathcal{N}(U \cap Y).$$

Meanwhile, we can also **restrict** an $\mathcal{O}_X$ module $\mathcal{M}$ to Y, because V being (affine) open in Y implies that it is (affine) open in X. This means that $\mathcal{M}(V)$ is defined. For example, restricting $\mathcal{O}_X$ to some open Y will give us the structure sheaf $\mathcal{O}_Y$ on Y – this is different from the direct image. If $\mathcal{M}$ is an $\mathcal{O}_X$ module, we'll say that the restriction of $\mathcal{M} = \mathcal{M}_X$ to an open Y is denoted $\mathcal{M}_Y$.

Example 172

What happens when we combine a restriction and a direct image: what is $j_* \mathcal{O}_Y$?

We know that the sections on $\mathcal{U}$ are the same as the sections of $\mathcal{O}_Y$ on $U \cap Y$ (which is a smaller open set in U), and this is the same as the regular functions on $U \cap Y$ – thus the sections can have poles in U but outside $U \cap Y$, and we have a map $\mathcal{O}_X \to j_* \mathcal{O}_Y$.

Our next topic is that of **twisting modules** on projective space $\mathbb{P}^r$. Recall that a **homogeneous fraction** $f = \frac{g}{h}$ is a fraction where g and h are each homogeneous: f has **degree** $\deg f = (\deg g - \deg h)$, and if g, h are relatively prime, then f is regular on an open set U if h does not vanish on U.

Definition 173

The sections of the **twisting module** $\mathcal{O}(n)$ on an open set $U \in \mathbb{P}^r$ are defined to be the homogeneous functions of degree n that are regular on U.

These are particularly important – we'll see why soon. Notably, if f is a homogeneous polynomial of degree d, then multiplying by f defines a homomorphism from $\mathcal{O}(n)$ to $\mathcal{O}(n+d)$, because f is regular everywhere.

Suppose that we restrict $\mathcal{O}(n)$ to the standard open set U^0 (that is, the points where $x_0 \neq 0$). Then the restriction is a **free $\mathcal{O}_{U^0}$-module** of rank 1 with basis $\{x_0^n\}$. This is because the homogeneous polynomials that are regular on all of U^0 are (scalar multiples of) powers of x_0 (everything else vanishes somewhere). Thus, the regular homogeneous fractions look like $f = \frac{g}{x_0^k}$ for some k – we can show the rest as an exercise.

Here's another definition of the twisting module: let H be the hyperplane where $x_0 = 0$. Then we can define the sections of $\mathcal{O}(nH)$ on U to be the rational functions f such that $f x_0^n$ is regular on U. So now f can have poles of order at most n on $U \cap H$. Now we have an isomorphism

$$\mathcal{O}(nH) \overset{x_0^n}{\to} \mathcal{O}(n).$$

We know that $\mathcal{O} \subset \mathcal{O}(H) \subset \mathcal{O}(2H) \cdots$, and we can consider

$$\lim_{n \to \infty} \mathcal{O}(nH).$$

This is the set of rational functions such that fx_0^n is regular for large enough n. Note that if fx_0^n is regular on an open set V, then f is a section of $\mathcal{O}$ on $V \cap U^0$, so f is an element of the direct image $j_*\mathcal{O}_{U^0}$. In other words,

$$\lim \mathcal{O}(nH) = j_*\mathcal{O}_{U^0},$$

where j is the inclusion of U^0 into $\mathbb{P}^r$.

We'll finish by discussing how to generate an $\mathcal{O}$-module $\mathcal{M}$. Suppose we have global sections $m = (m_1, \cdots, m_k)$ of $\mathcal{M}$ – they are elements of $\mathcal{M}(\mathbb{P}^r)$, and we get homomorphisms $O^k \overset{m}{\to} \mathcal{M}$, sending

$$(\alpha_1, \cdots, \alpha_k) \to \sum \alpha_i m_i.$$

When this map is surjective, we say that m **generates** the $\mathcal{O}_{\mathbb{P}^r}$ module $\mathcal{M}$. Then $\mathcal{M}$ is a finite $\mathcal{O}$-module for every affine open set $U = \operatorname{Spec} A$ if $M = \mathcal{M}(U)$ is a finite $A = \mathcal{O}(U)$-module. (Note that this is only true because U is affine!)

We'll introduce the notation

$$\mathcal{M}(n) = \mathcal{M} \otimes_{\mathcal{O}} \mathcal{O}(n).$$

Theorem 174

Say that $\mathcal{M}$ is a finite $\mathcal{O}$-module on $\mathbb{P}^r$. Then the twisting module $\mathcal{M}(n)$ can be (finitely) generated by global sections for sufficiently large n.

$\mathcal{M}$ itself is not generated by global sections – this is only true on the affines, and $\mathcal{M}$ does not need to have any sections at all. We'll review this a bit at the beginning of next lecture and then move on to cohomology.

24 April 13, 2020

Our goal for today is to show the result from the end of last time: given a finite $\mathcal{O}$-module $\mathcal{M}$ on $\mathbb{P}^n$, we can generate the twisting module $\mathcal{M}(n)$ by global sections. Recall that the **twisting module** $\mathcal{O}(n)$ is defined such that the sections on an open set U are the homogeneous fractions $f = \frac{g}{h}$ are regular on all of U. For the hyperplane H at infinity, we can define $\mathcal{O}(nH)$ similarly: its sections on U are the homogeneous rational functions of degree 0 such that $x^n\alpha$ is regular of degree n. Then multiplying by x^n gives us an isomorphism from $O(nH)$ to $O(n)$. Recall that $\mathcal{O} \subset \mathcal{O}(H) \subset \mathcal{O}(2H) \cdots$ and so on, and taking the limit as $n \to \infty$ yields a limit $\lim \mathcal{O}(nH)$ such that its sections on an open set U are the rational functions α with $x_0^n\alpha$ regular for sufficiently large n.

We discussed that **this limit** $\lim \mathcal{O}(nH)$ **is the direct image** $j_*\mathcal{O}_{U^0}$, where j is the inclusion of U^0 into $X = \mathbb{P}^r$. Recall that the sections of the direct image are defined as

$$[j_*\mathcal{O}_U^0](V) = \mathcal{O}_{U^0}(V) = O_X(V),$$

which is the set of regular functions on V. Let's show this explicitly:

Proof. A homogeneous fraction $f = \frac{g}{h}$ (for relatively prime g, h) is regular on an open set U^0 if the denominator doesn't vanish on U^0, meaning it must be a power of x_0 (times a scalar). To show when f is regular on an open set $V \cap U^0$, we factor the denominator $h = h_1 \cdots h_k$, where the h_i are irreducible polynomials. Then we care about having $h \neq 0$ everywhere on $V \cap U^0$, and this happens whenever all $h_i \neq 0$ on $V \cap U^0$. Letting Y_i be the set of zeros of h_i, we have an **irreducible** closed set in $X = \mathbb{P}^r$. But $V \cap U^0$ is open, so $Y_i \cap (V \cap U^0)$ is open in Y_i, meaning this set is **dense** or empty.

And now if $h_i \neq 0$, we're in the case where we want $Y_i \cap (V \cap U^0)$ must be empty. We can write this as $(Y_i \cap V) \cap (Y_i \cap U^0)$, and both terms here are either dense in Y_i or empty – thus we need at least one of them to be empty. In the latter case, we have that $h_i = cx_0$ (analogous to the simple case of being regular on U^0), which we've analyzed before. But in the former case, we have $h_i \neq 0$ on all of V. Thus, $f = \frac{g}{h}$ is regular on $V \cap U^0$ if $h = h' x_0^k$, where h' is a regular function on V. (Basically, we have a power of x_0, and the rest can't vanish on all of V.) Therefore we need that $f x_0^k$ is regular on V if k is sufficiently large.

But now $[j_* \mathcal{O}_{U^0}](V) = \mathcal{O}_{U^0}(V \cap U^0)$ is the set of rational functions of degree 0 satisfying the condition that $x_0^k \alpha$ is regular on V for sufficiently large k: thus we do have $j_* \mathcal{O}_{U^0}$ as the limit $O(nH)$. □

Remember that we can **twist a module**: because $\mathcal{O} \subset \mathcal{O}(H) \subset \cdots$, we can define

$$\mathcal{M}(n) = \mathcal{M} \otimes_{\mathcal{O}} \mathcal{O}(n), \quad \mathcal{M}(nH) = \mathcal{M}M \otimes \mathcal{O}(nH).$$

(These maps may not be injective if $\mathcal{M}$ has torsion.) Then the limit here is

$$\boxed{\lim_{n \to \infty} \mathcal{M}(nH)} = \mathcal{M} \otimes_{\mathcal{O}} \lim \mathcal{O}(nH) = \mathcal{M} \otimes_{\mathcal{O}} j_* \mathcal{O}_{U^0} = \boxed{j_* \mathcal{M}_{U^0}}.$$

This is important, and we'll discuss it a bit more: if we're looking at a standard affine open of $\mathbb{P}^r$, then $\mathcal{O}(U^i)$ are the rational functions with denominator x_i^k. Then the sections of $\mathcal{O}(n)$ on U^i are homogeneous fractions $f = \frac{g}{x_0^k}$ of degree n, so g is of degree $n + k$. Then $[\mathcal{O}(nH)](U^i)$ is the set of fractions of degree 0 such that multiplying by x_0^n yields a regular function on U^i, so $x_0^n f = \frac{g}{x_i^k} \implies f = \frac{g}{x_0^m x_i^k}$, where k is aribtrary, $m \leq n$, and g is a function of degree $(n + k)$.

But now $\mathcal{O}$ maps isomorphically to $\mathcal{O}(n)$ on U^i by the multiply-by-x_i^n map, and we have an isomorphism on X from $\mathcal{O}(n)$ to $\mathcal{O}(nH)$ by dividing by x_0^n. So then an $\mathcal{O}$-module $\mathcal{M}$ goes to $\mathcal{M}(n)$ via the isomorphism map $1 \otimes x_i^n$, while $\mathcal{M}(n)$ goes to $\mathcal{M}(nH)$ via the isomorphism map $1 \otimes x_0^{-n}$.

So now let's go back to the result we're trying to show: recall that a $\mathcal{O}$-module $\mathcal{M}$ is a **finite $\mathcal{O}$-module** if for every affine open set $U = \operatorname{Spec} A$, $\mathcal{M}(U)$ is a finite $\mathcal{O}(U)$-module (that is, a finite A-module). We discussed trying to generate an $\mathcal{O}$-module $\mathcal{M}$ by **global sections** $m_1, \cdots, m_k \in \mathcal{M}(X)$ by considering the homomorphism

$$(\alpha_1, \cdots, \alpha_k \to \sum \alpha_i m_i,$$

and the homomorphism being surjective means that we **generate** $\mathcal{M}(X)$.

Proof of theorem. Let $M = \mathcal{M}(U^0)$ and $A = \mathcal{O}(U^0)$. M is a finite A-module, and it is also $[\mathcal{M}(U)](U^0)$.

It's enough to show that these global sections of $\mathcal{M}(nH)$ restrict to generators of A mod M. This is because we know that whenever $U = \operatorname{Spec} A$ is affine, $\mathcal{O}_U$-modules M correspond to A-modules M. And we know how to localize this, so we can extend to the localizations of U, which is enough. Then if the global sections of $\mathcal{M}(nH)$ generate M, then they generate $\mathcal{M}_{U^0}$. We have a map $\mathcal{O}_{U^0}^k \to \mathcal{M}_{U^0}$ – let C be the cokernel of the supposed generators. We know that the cokernel C_{U^0} is zero here, and we can similarly show that C_{U^i} is zero (possibly with a larger value of n), and then we use the sheaf axiom: remember that we have an injective map which can be represented as $0 \to C(V) \to \prod_i C(V^i)$.

But $\lim \mathcal{M}(nH) = j_* \mathcal{M}_{U^0}$, where j is again the inclusion map, and we can pick generators for $M = \mathcal{M}(U^0)$ – this is supposed to be a finite A-module. We can represent by global sections of $\mathcal{M}(nH)$ – call them $m_1', \cdots, m_k'$ – since these map to $m_1, \cdots, m_k$, these generate M. □

25 April 15, 2020

We'll begin talking about **cohomology** of $\mathcal{O}$-modules today. There isn't really any easy way to uniquely describe what's going on here, though there are lots of different constructions.

Let X be a variety, and consider the **short exact sequence**

$$0 \to \mathcal{L} \xrightarrow{f} \mathcal{M} \xrightarrow{g} \mathcal{N} = 0.$$

This means that f is injective, $\ker g = \operatorname{Im} f$, and g is surjective. Then if $U = \operatorname{Spec} A$ is an affine open set of X, we can map the sections via

$$0 \to \mathcal{L}(U) \overset{f(U)}{\to} \mathcal{M}(U) \overset{g(U)}{\to} \mathcal{N}(U) \to 0 :$$

this is an exact sequence of A-modules, where $A = \mathcal{O}(U)$. However, the sequence may not be exact if U is open but not completely affine:

Lemma 175

If we have the short exact sequence $0 \to \mathcal{L} \to \mathcal{M} \to \mathcal{N} \to 0$ above, then $0 \to \mathcal{L}(U) \overset{f(U)}{\to} \mathcal{M}(U) \overset{g(U)}{\to} \mathcal{N}(U)$ is exact for any open set U. (We do not need the surjectivity of g for this.)

Proof. Remember that an $\mathcal{O}$-module is defined on the affine opens, and it extends to all opens by the sheaf property of $\mathcal{L}, \mathcal{M}$, and $\mathcal{N}$. If $\{U^i\}$ is an affine cover of U, then we have the exact sequence

$$0 \to \mathcal{M}(U) \to \prod_i \mathcal{M}(U^i) \to \prod_{i,j} \mathcal{M}(U^i \cap U^j)$$

by the definition of $\mathcal{M}(U)$. We can then draw a commutative diagram:

$$\begin{tikzcd}
& 0 \arrow[d] & 0 \arrow[d] & 0 \arrow[d] \\
0 \arrow[r] & \mathcal{L}(U) \arrow[r] \arrow[d] & \mathcal{M}(U) \arrow[r] \arrow[d] & \mathcal{N}(U) \arrow[d] \\
0 \arrow[r] & \prod_i \mathcal{L}(U^i) \arrow[r] \arrow[d] & \prod_i \mathcal{M}(U^i) \arrow[r] \arrow[d] & \prod_i \mathcal{N}(U^i) \arrow[d] \\
0 \arrow[r] & \prod_{i,j} \mathcal{L}(U^{ij}) \arrow[r] & \prod_{i,j} \mathcal{M}(U^{ij}) \arrow[r] & \prod_{i,j} \mathcal{N}(U^{ij})
\end{tikzcd}$$

The columns are exact by the sheaf property, and because the U^is and U^{ij}s are affine, the last two rows are exact by definition (of $0 \to \mathcal{L} \to \mathcal{M} \to \mathcal{N} \to 0$ being exact). Our goal is to prove that the first row is also exact (in other words, that the additional blue 0 and arrows are valid).

Take an element of $\mathcal{L}(U)$ that goes to zero in $\mathcal{M}(U)$, and take its image in $\mathcal{L}(U^i)$. $\prod \mathcal{L}(U^i)$ is injective into $\prod_i \mathcal{M}(U^i)$, so it goes to 0, but then the $\mathcal{L}$ column is zero, so $\mathcal{L}(U)$ must be zero. We can use a similar Snake Lemma idea to show the middle part works out too. □

Basically, the sections are **left exact** but not necessarily right exact.

The next idea is to study **global sections**: by the above lemma, we know that we have the exact sequence

$$0 \to \mathcal{L}(X) \to \mathcal{M}(X) \to \mathcal{N}(X).$$

Then the **cohomology** is a sequence of functors $H^0, H^1, \cdots$, from $\mathcal{O}$-modules to vector spaces, which "substitutes for the lack of exactness."

This has three characteristic properties:

- $H^0(\mathcal{M}) = \mathcal{M}(X)$ are just the global sections.
- We have the **cohomology sequence** which describes the lack of exactness: if we have a short exact sequence $0 \to \mathcal{L} \to \mathcal{M} \to \mathcal{N} \to 0$, we get the long exact sequence

$$0 \to H^0(\mathcal{L}) \to H^0(\mathcal{M}) \to H^0(\mathcal{N}) \xrightarrow{\delta^0}$$

$$\to H^1(\mathcal{L}) \to H^1(\mathcal{M}) \to H^1(\mathcal{N}) \xrightarrow{\delta^1} \cdots$$

  δ^0 is called the **coboundary map**, so computing $H^1(\mathcal{L})$ will tells us to what extent exactness fails. We'll see later that all but finitely many of these are zero.

- Basically, if X is affine, then the global section functor is exact, so we don't need higher cohomology. The actual statement is a bit more complicated: if $U = \operatorname{Spec} A$ is an affine open subset of X, and $\mathcal{N}$ is an $\mathcal{O}_U$-module, then $H^q(j_*\mathcal{N}) = 0$ for all **positive** q. (Meanwhile, $H^0(j_*\mathcal{N})$ is defined to be $(j_*\mathcal{N})(X) = \mathcal{N}(U)$, so it is not zero.)

Corollary 176

If X is affine, then $H^q(\mathcal{M}) = 0$ for all $q > 0$ and all $\mathcal{M}$.

Proof. Use the identity map $X \to X$ with the third charactersitic property. Then $j_*\mathcal{N}$ is just $\mathcal{N}$. □

Theorem 177

There is a cohomology theory satisfying the three characteristic properties above, unique up to isomorphism.

Unfortunately, there is no natural construction of the cohomology – the best way to think about it is to just work with the characteristic properties.

Proof of uniqueness of cohomology. Say we are given a cohomology. Choose an affine open cover $U = \{U^\nu\}$ of X – it will follow from the logic that the choice of affine cover doesn't matter. Let j be the family of inclusions of U into X. Let $\mathcal{M}$ be an $\mathcal{O}$-module – we want to describe its cohomology. We can restrict to open subsets: let

$$\mathcal{M}_U = \prod \mathcal{M}_{U^\nu}, \quad \prod j_*^\nu \mathcal{M}_{U^\nu}.$$

Lemma 178

Let $\mathcal{R} = j_*\mathcal{M}_\nu$. There exists a canonical injective map from $\mathcal{M} \to \mathcal{R}$, and $H^q(\mathcal{R}) = 0$ if $q > 0$.

Proof. Let V be open in X. We know that $V^\nu = V \cap U^\nu$ cover V, so $\mathcal{M}(V)$ is contained in the product $\prod M(V^\nu)$, which we can write as a product $\prod M_{U^\nu}(V)$. By definition, this is $R(V)$, so $\mathcal{M}(V)$ is contained in $\mathcal{R}(V)$, meaning we do have injectivity.

For the second part,

$$H^q(\mathcal{R}) = \prod H^q(j^\mu_* \mathcal{M}_{U^\nu}) = 0,$$

because each U^ν is affine. □

Now we can form a short exact sequence $0 \to \mathcal{M} \to \mathcal{R} \to \mathcal{S} \to 0$, where $\mathcal{S}$ is the cokernel of the map from $\mathcal{M}$ to $\mathcal{R}$. We know that the correspondending cohomology sequence has

$$0 \to H^0(\mathcal{M}) \to H^0(\mathcal{R}) \to H^0(\mathcal{S}) \to$$

$$H_1(\mathcal{M}) \to 0 \to H^1(\mathcal{S}) \ldots,$$

but the key point here is that in the exact sequence

$$\boxed{0 \to H^0(\mathcal{M}) \to H^0(\mathcal{R}) \to H^0(\mathcal{S}) \overset{\delta^0}{\to} H^1(\mathcal{M})}$$

with the cohomology given and a chosen covering of X, we know what the H^0s look like – this gives us the first three terms of the exact sequence, because $H^0(\mathcal{R})$ depends on the covering rather than the cohomology.

From here, we can look at $0 \to H^1(\mathcal{S}) \overset{d}{\to} H^2(\mathcal{M}) \to 0$: $H^1(\mathcal{S})$ is determined by $\mathcal{M}$, not by the cohomology, so $H^1(\mathcal{S})$ is unique, and therefore $H^2(\mathcal{M})$ is unique. Now we can repeat this argument inductively to show that everything is unique: we do need to check uniquness in the δs, but the boxed sequence identifies $H^1(\mathcal{M})$ uniquely independent of the cohomology – this fixes δ, and the idea is that knowing the coboundary maps is enough. □

26 April 17, 2020

Recall the characteristic propertes of cohomology, which is a series of functors $H^0, H^1, H^2, \cdots$. Here, $\mathcal{M}$ is an $\mathcal{O}_X$-module for some variety X.

- $H^0(X, \mathcal{M}) = H^0(\mathcal{M}) = \mathcal{X}$ is just the vector space of global sections. But there isn't a very clear correspondence: it's not clear that when $\mathcal{M}$ is a finite module over the coordinate ring, $H^0(\mathcal{M})$ forms a vector space over the scalars.
- We have the cohomology sequence: if we have a short exact sequence $0 \to \mathcal{L} \to \mathcal{M} \to \mathcal{N} \to 0$, this yields a sequence of maps $H^q(\mathcal{L}) \to H^q(\mathcal{M}) \to H^q(\mathcal{M})$ (because H^q is a functor, these maps exist by definition), and then we have a coboundary map δ to $H^{q+1}(\mathcal{L})$ such that the entire sequence is exact.
- If U is an affine open set in X, j is the inclusion map $U \to X$, and $\mathcal{N}$ is an $\mathcal{O}_U$-module, then $H^q(X, j_*\mathcal{N}) = 0)$ for all $q > 0$. In other words, we don't really have a cohomology for affine open sets – it just looks the same as the exact sequence we already have.

We can talk about the third point here in more generality:

Lemma 179

If Y is affine, $f : Y \to X$ is a morphism, and $\mathcal{N}$ is an $\mathcal{O}_Y$-module, then $H^q(X, f_*\mathcal{N}) = 0$ for all $q > 0$.

This has consequences for what are called **affine morphisms**:

Definition 180

Let Y and X be varieties. A morphism $f : Y \to X$ is an **affine morphism** if for all affine open sets U in X, $V = f^{-1}(U)$ is affine in Y.

Example 181

The most important case is when f is an inclusion of a closed (or open) subvariety Y into X: this is indeed an affine open set.

Theorem 182

Let $f : Y \to X$ be an affine morphism, and let $\mathcal{N}$ be an $\mathcal{O}_Y$-module. Then the cohomology

$$H^q(Y, \mathcal{N}) \cong H^q(X, f_*\mathcal{N}).$$

Proof. We want to show that the sequence of $H^q(X, f_*\mathcal{N})$s is the cohomology for $\mathcal{N}$, so we just need to verify the three characteristic properties. Define $F^q(\mathcal{N}) = H^q(X, f_*\mathcal{N})$ for convenience.

For the first characteristic property, we wish to show that $F^0(\mathcal{N}) = \mathcal{N}(Y)$ forms the space of global sections. This is $H^0(X, f_*\mathcal{N})$, which is the space of global sections of $f_*\mathcal{N}$ on X. But f_* is defined to have the sections of $\mathcal{N}$ on the inverse image of X, which means this is indeed $\mathcal{N}(Y) = H^0(Y, \mathcal{N})$.

For the second property, we need to look at a short exact sequence $0 \to \mathcal{L} \to \mathcal{M} \to \mathcal{N} \to 0$ of $\mathcal{O}_Y$-modules. Because f is an affine morphism, we claim that the sequence $0 \to f_*\mathcal{L} \to f_*\mathcal{M} \to f_*\mathcal{N} \to 0$ is also exact. **But the definition of being exact depends on being exact on the affine open sets**, and if U is open in X, we want to know whether the sequence

$$0 \to f_*\mathcal{L}(U) \to f_*\mathcal{M}(U) \to f_*\mathcal{N}(U) \to 0$$

is also exact. But letting $V \in Y$ be the inverse image $f^{-1}(U)$, we konw that $f_*\mathcal{L}(U) = \mathcal{L}(V)$, so that sequence is actually just

$$0 \to \mathcal{L}(V) \to \mathcal{M}(V) \to \mathcal{N}(V) \to 0,$$

and this is exact because V is an affine open set of Y and $0 \to \mathcal{L} \to \mathcal{M} \to \mathcal{N} \to 0$ is exact.

For the third property, we take an affine open $V \in Y$ and an $\mathcal{O}_V$-module $\mathcal{N}$, and we let $j : V \to Y$ be the inclusion map. We want to show that $F^q(j_*\mathcal{N}) = 0$ for all $q > 0$, and this is just

$$F^q(j_*\mathcal{N}) = H^q(X, f_*j_*\mathcal{N}).$$

The composition $f_* \circ j_*$ is a morphism, and V is affine, so Lemma 179 tells us that this is zero, as desired. □

Example 183

Our goal with this will be to compute a specific cohomology: that of the twisting module $\mathcal{O}(n)$ on projective space $\mathbb{P}^d$.

First of all, $H^0(\mathbb{P}^d, \mathcal{O}(n)$ is the set of homogeneous fractions of degree n, but being a global section means that our denominator must be a constant (so it doesn't vanish anywhere), so this is the space of homogeneous polynomials of degree n, which is a $\binom{n+d}{d}$-degree vector space whenever $n \ge 0$.

To proceed, let's map $\mathcal{O}_\mathbb{P}(n) \to \mathcal{O}_H(n)$ to the hyperplane at infinity where $x_0 = 0$: we just send x_0 to 0. The kernel is the set of homogeneous polynomials of degree $n-1$ times x_0, so we get the exact sequence

$$0 \to \mathcal{O}_\mathbb{P}(n-1) \xrightarrow{\cdot x_0} \mathcal{O}_\mathbb{P}(n) \to \mathcal{O}_H(n) \to 0.$$

This gives us a cohomology sequence: we'll just write the dimensions of the spaces instead of the spaces themselves, which yields

$$0 \to \binom{(n-1)+d}{d} \to \binom{n+d}{d} \to \binom{n+(d-1)}{d-1}.$$

(Indeed, for this to be an exact sequence, the dimension of the middle space is supposed to be the sum of the dimensions of the two outside spaces, and this is Pascal's identity.)

Theorem 184

We have $H^q(\mathcal{O}(n)) = 0$ for all $q > 0$ and $n \ge 0$.

Proof. We induct on d: we can assume this is true for the cohomology on H, which is $\mathbb{P}^{d-1}$. So then we can assume the third columns of the dimensions of H^q are all zero. But the first columns all have dimension zero by induction on n, so the second column all has dimension zero. □

Theorem 185

Suppose $n = -r$ for some $r > 0$. Then the dimension $h^q(\mathcal{O}(-r)) = 0$ if $q \neq d$, and the dimension of the dth cohomology

$$h^d(\mathcal{O}(-r)) = \binom{r-1}{d}.$$

Proof. First, we do $r = 1$. We look at our exact sequence (where i_* is the inclusion of H into $\mathbb{P}$)

$$0 \to \mathcal{O}_\mathbb{P}(-1) \to \mathcal{O}_\mathbb{P} \to i_*\mathcal{O}_H \to 0:$$

the dimensions on the seccond and third columsn are 1 on the first cohomology (corresponding to the constant functions, the scalars), and zero everywhere else. So all of the cohomology on the first column$\mathcal{O}_\mathbb{P}(-1)$ is 0 for $\mathcal{O}(-1)$. And then an induction argument works just like for the positive case: consider

$$0 \to \mathcal{O}_\mathbb{P}(-r-1) \to \mathcal{O}_\mathbb{P}(-r) \to \mathcal{O}_\mathbb{H}(-r) \to 0,$$

inducting on r in the second column and d in the third. □

Example 186

Let Y be a plane curve of degree r in $\mathbb{P}^2$. Then we have (where i_* is the inclusion $Y \to \mathbb{P}^2$)

$$0 \to \mathcal{O}_\mathbb{P}(-r) \xrightarrow{\cdot f} \mathcal{O}_\mathbb{P} \to i_*\mathcal{O}_Y \to 0.$$

(a regular function on $\mathbb{P}^2$ restricts to a regular function on Y, and the functions that are 0 on Y are those divisible by f).

We've now computed the dimensions in the first two columns: column 1 ($\mathcal{O}_{\mathbb{P}}(-r)$ has $0, 0, \binom{r-1}{2}$, followed by a bunch of zeros, and column 2 ($\mathcal{O}_{\mathbb{P}}$) has 1 followed by a bunch of 0s. So then using properties of exactness gives us that column 3 ($\mathcal{O}_Y$) has dimensions $1, \binom{r-1}{2}, 0, \cdots$. In other words,

$$h^1(Y, \mathcal{O}_Y) = p_a$$

is the **arithmetic genus** of Y, which is also equal to g, the **geometric genus**, when Y is smooth.

In both of the above cases, it was okay to just ignore the i_*, because the cohomology transfers over:

$$H^q(Y, \mathcal{O}_Y) \cong H^q(\mathbb{P}^2, i_*\mathcal{O}_Y),$$

because i_* is an affine morphism.

Theorem 187

Let $\mathcal{M}$ be a finite $\mathcal{O}_X$-module, where $X = \mathbb{P}^n$. Then

$$H^q(\mathcal{M}(n)) = 0$$

for sufficiently large "twistings" n.

Proof. We know that $\mathcal{M}(r)$ is generated by global sections for sufficiently large r: thus, there exists a map from $\mathcal{O}^m \to \mathcal{M}(r)$ (where $\mathcal{O}$ is the structure sheaf – sufficiently large twists are finitely generated), and if K is the kernel, we have the exact sequence

$$0 \to K \to \mathcal{O}^m \to \mathcal{M}(r) \to 0.$$

Twist this sequence, which yields

$$0 \to K(r) \to \mathcal{O}(n)^m \to \mathcal{M}(n+r) \to 0.$$

We know that the middle column has dimension 0 for $q > 0$ and $n \ge 0$ by Theorem 184. But then we have

$$H^q(\mathcal{O}(n)^m) \to H^q(\mathcal{M}(n+r)) \xrightarrow{\delta} H^{q+1}(K(n)) \to H^{q+1}(\mathcal{O}(n)^m),$$

where the first and last spaces are 0.

Lemma 188

$H^q(\mathcal{M}) = 0$ for sufficiently large q.

Apply this to the coboundary map, which is an isomorphism: $H^q(\mathcal{M}(n+1)) \to H^{q+1}(K(n))$ for sufficiently large q, so because $H^{q+1}(K(n))$ is zero for sufficiently large q, $H^q(\mathcal{M}(n+r))$ is also zero for sufficiently large q. □

27 April 22, 2020

Due to popular vote, we'll be replacing the quiz with a final assignment: it will be more cumulative but also more routine than the usual homework questions.

We'll begin by discussing **support**:

Definition 189

Suppose that A is a finite type domain for simplicity and M is a finite A-module. The **annihilator** of M, denoted $\text{ann } M$, is an ideal of A: it's the set of α such that $\alpha M = 0$. The **support** of M, denoted $\text{supp } M$, is the zero locus $V(\text{ann } M)$ of the annihilator in $\operatorname{Spec} A$.

This definition carries over to $\mathcal{O}$-modules, because both the support and annihilator are compatible with localization.

Proposition 190

Let $\overline{M} = M/(\mathfrak{m}_p M) = M \otimes_A k(p)$. Then

$$\text{supp } M = \{p \in X : \overline{M} \neq 0\}.$$

In other words, the support of M is the "set of points of M that are not zero."

Proof. Let I be the annihilator $\text{ann } M$. If $I \not\subset \mathfrak{m}_p$, then there exists $\alpha \in I$ such that $\alpha(p) \neq 0$. But on $\overline{M}$, scalar multiplication by α is the same as multiplying by $\alpha(p)$. If this value is nonzero, then $\alpha\overline{M} = \alpha(p)\overline{M} = \overline{M}$ is not zero.

On the other hand, if $\overline{M} = 0$, then $\mathfrak{m}_p M = M$, and now Nakayama tells us that there exists an element $z \in \mathfrak{m}_p$ such that $(1 - z)M = 0$. Thus an element of I, $1 - z \in \text{ann } M$, so I is not contained in the maximal ideal. □

The next result we'll cover today is the **vanishing of cohomology for large dimension**. First of all, suppose that X is a projective variety and we have an inclusion $\iota : X \to \mathbb{P}$. Let $\mathcal{M}$ be an $\mathcal{O}_X$-module: we know that

$$H^q(X, \mathcal{M}) \cong H^q(R, \iota_*\mathcal{M}),$$

because ι is an affine morphism. So we can work with cohomology on projective space instead, which is convenient because we know how twisting works there.

So now we can let $\mathcal{M}$ be an $\mathcal{O}$-module on $\mathbb{P}$, and let $U^0 : x_0 \neq 0$ be the standard affine open set. If we let $j : U^0 \to \mathbb{P}$ be the inclusion map, we have the sequence

$$\mathcal{M} \to \mathcal{M}(H) \to \mathcal{M}(2H) \to \cdots,$$

where $\mathcal{M}(H) = \mathcal{M} \otimes_{\mathcal{O}} \mathcal{O}(H)$, such that the limit $\mathcal{M}(nH)$ is the direct image $j_*\mathcal{M}_{U^0}$.

Theorem 191

Let $\mathcal{M}$ be a finite $\mathcal{O}$-module on projective space, and suppose $\text{supp } m$ has dimension k. Then $H^q(\mathcal{M}) = 0$ for all $q > k$.

Proof. Consider the map from $\mathcal{M}(-1)$ to $\mathcal{M}$, given by multiplying by x_0. Embed this in the exact sequence

$$0 \to K \to \mathcal{M}(-1) \to \mathcal{M} \to C \to 0.$$

The twists have the same support as $\mathcal{M}$: $\mathcal{O}(n)$ is locally isomorphic to $\mathcal{O}$. x_0 annihilates K and C, so we can assume that the support of K (some closed set in projective space) is not contained in the hyperplane, meaning that the dimensions of the support of K and C are both less than k. (The support of K is contained in the support of $\mathcal{M}$, and we can assume that x_0 is not zero on any component of the support of $\mathcal{M}$.)

Here, we have the exact sequences

$$0 \to \mathcal{K} \to \mathcal{M}^{-1} \to \mathcal{M} \to 0$$

which corresponds to the cohomology sequence

$$H^q(\mathcal{K}) \to H(\mathcal{M}(-1)) \to H^q(\mathcal{M}) \to \mathcal{H}^{q+1}.$$

By induction, we know that the first and last of these are zero, so we have an isomorphism between $H^q(\mathcal{N})$ and $H(\mathcal{M}(-1))$. Similarly,

$$0 \to \mathcal{M}(-1) \to \mathcal{M} \to C \to 0$$

means that $H^q(\mathcal{M}(-1))$ is isomorphic to $\mathcal{H}^q(\mathcal{M})$.

Now we can twist: replacing $\mathcal{M}$ by $\mathcal{M}(n)$,

$$H^q(\mathcal{M}(n-1)) \cong H^q(\mathcal{M}(n))$$

for $n > k$. And then taking the limit, for all $q > k$,

$$\lim H^q\mathcal{M}(n) = H^q(\lim \mathcal{M}(n)) = H^q(j_*\mathcal{M}^0_U) = 0,$$

because U^0 is an affine open set. This means that H^q must be zero for sufficiently large 0. □

Corollary 192

If $\mathbb{P} = \mathbb{P}^n$, then $H^q(\mathcal{M}) = 0$ for all $q > n$.

Our next result basically says that "twisting kills the cohomology:"

Theorem 193

Let $\mathcal{M}$ be a finite $\mathcal{O}$-module on $\mathbb{P}$. Then $H^q(\mathcal{M}(n)) = 0$ for all $q > 0$ and n sufficiently large.

Proof. Recall that $\mathcal{M}(r)$ is generated by global sections for large enough r, which means we have a surjective map represented by the exact sequence $\mathcal{O}^k \to \mathcal{M}(r) \to 0$. Letting $\mathcal{N}$ be the kernel, we have some exact sequence

$$0 \to \mathcal{N} \to \mathcal{O}^k \to \mathcal{M}(r) \to 0.$$

Twist this sequence to get

$$0 \to \mathcal{N}(d) \to \mathcal{O}(d)^k \to \mathcal{M}(d+r) \to 0,$$

and we know this yield the cohomology map

$$H^q\mathcal{O}(d)^k \to H^q\mathcal{M}(d+r) \to H^{q+1}\mathcal{N}(d) \to H^{q+1}\mathcal{O}(d)^k,$$

and the first and last thing here are zero, so the middle two are isomorphic. So by the previous result, we know that $H^q(\mathcal{M}(d+r)) = 0$ for large enough q, and the same logic shows that $H^q\mathcal{N}(d') = 0$ for some large enough d'. Thus $H^{q-1}\mathcal{M}(d'+r) = 0$, and we can keep repeating this down for smaller and smaller q (adding some r each time). □

Theorem 194

Let $\mathcal{M}$ be a finite module on $\mathbb{P}$. Then $H^q(\mathcal{M})$ is a finite-dimensional vector space for all q.

Proof for $q > 0$. As above, embed the multiply-by-x_0-map in an exact sequence

$$0 \to K \to \mathcal{M}(-1) \to \mathcal{M} \to C \to 0.$$

We can induct on the dimension of the support: because K and C are of lower dimension, it suffices to assume H^qK, H^qC are finite-dimensional vector spaces. We have the exact sequences

$$H^qK \to H^q(\mathcal{M})(-1) \to H^q\mathcal{M} \to H^{q+1}K,$$

where the first and last term are both finite-dimensional vector spaces. Then the middle two must be both finite-dimensional or both infinite-dimensional. The same argument works for the sequence

$$H^{q-1}C \to H^{q-1}\mathcal{M} \to H^q\mathcal{M} \to H^qC.$$

so we know that $H^q\mathcal{M}(-1)$ and $H^q(\mathcal{M})$ are either both finite-dimensional or both infinite-dimensional. Now we can twist and use the previous theorem: $H^q(\mathcal{M}(n))$ will eventually have cohomology zero, and we can use descending induction on n to show that all the twists have finite dimension. □

The $q = 0$ case is harder to prove directly – it's not obvious that the global sections form a finite-dimensional vector space.

28 April 24, 2020

We'll discuss the cohomology of projective smooth curves today – unfortunately, we won't get time to talk about surfaces.

First of all, a review: we define cohomology on projective space $\mathbb{P}$ for a finite $\mathcal{O}$-module (finiteness isn't officially necessary, but it makes things easier for our discussion – for example, the support exists). There are three theorems we shuld know about:

- If the dimension of the support supp $\mathcal{M}$ is k, then $H^q\mathcal{M} = 0$ for all $q > k$.
- The cohomology of the twist $H^q(\mathcal{M}M(n))$ is 0 for sufficiently large n for any $q > 0$. (This one will be less important today.)
- $H^q(\mathcal{M})$ is a finite-dimensional vector space (this requires finite module).

These come from a paper by Serre from 1955 – this was also the first time the concept of an $\mathcal{O}$-module was introduced.

One other fact we should know: if $i_* : X \to \mathbb{P}$ is an inclusion of a projective variety, and $\mathcal{M}$ is an $\mathcal{O}_X$-module, then

$$H^q(X, \mathcal{M}) \cong H^q(\mathbb{P}, i_*\mathcal{M}).$$

This means that both the first and third result that we've stated above always hold for any projective variety.

We'll now move to new materials. Recall that local rings of a smooth projective curve X are valuation rings: let v_p be the valuation at the point $p \in X$.

Definition 195

A **divisor** on a smooth projective curve X is a finite combination of points

$$D = \sum_i r_i p_i, \quad r_i \in \mathbb{Z}.$$

The **divisor of a rational function** f is

$$\text{div} f = \sum_{p \in X} v_p(f) p.$$

Since the valuation is zero on all but finitely many points, the divisor of a rational function is indeed a finite sum: **it's basically the sum of the zeros minus the sum of the poles**, counting multiplicity. Here's how we can make that more precise:

Definition 196

If $v_p(f) > 0$, then f has a **zero of order** $r = v_p(f)$. Meanwhile, when $v_p(f) < 0$, then f has a **pole of order** $r = -v_p(f)$.

Definition 197

A divisor D is **effective** if $r_i \ge 0$ for all i: then we say that $D \ge 0$. Given an open set U, D is **effective on U** if $r_i \ge 0$ for all $p_i \in U$.

One classical problem of algebraic geometry is to determine the rational functions f whose poles are bounded by a given effective divisor, and we'll use cohomology to discuss this.

Definition 198

An **$\mathcal{O}$-module associated to a divisor** D is defined such that its sections on $\mathcal{O}(D)$ on an open set U are the rational functions f such that div $f + D$ is effective, plus the zero function.

By definition, the cohomology $H^0(\mathcal{O}(D)$ is exactly this set of rational functions f such that div $f + D \ge 0$, as well as zero. This is the set that we're trying to understand in the classical problem.

Lemma 199

$\mathcal{O}(D)$ is an $\mathcal{O}$-module, and it is **locally free** of rank 1 (so it is free on each open set, meaning it looks like $\mathcal{O}$).

Example 200

Consider the divisor $D = -p$.

This is the set of points such that div $\text{f} - p \ge 0$, meaning that f must have a zero at p (so that it cancels out with the $-p$). Thus, on any open set U that contains p,

$$\mathcal{O}(-p)(U) = \text{ functions with a zero at } p = \mathfrak{m}_p.$$

On the other hand, if U does not contain p, then the sections are just the regular functions.

Example 201

Now consider the divisor $D = p$.

Then our functions are allowed to have a pole of order at most 1 at p, so on any open set U that contains p,

$$\mathcal{O}(p)(U) = \text{ rational functions with a possible pole at } p.$$

(And again, we just get the regular functions for any U not containing p.)

The reason we care about cohomology is that it tells us that **dimension** of these spaces. We'll start by looking at our first example with the maximal ideal: we have the exact sequence

$$0 \to \mathcal{O}(-p) \to \mathcal{O} \to \kappa_p \to 0,$$

where κ_p is the one-point module $\kappa(p)$. We'll tensor this with $\mathcal{O}(D)$: the tensor product is right exact normally, but because $\mathcal{O}(D)$ is locally free, we'll still have the whole exact sequence.

Lemma 202

The module associated to $\mathcal{O}(D + E)$ is isomorphic to $\mathcal{O}(D) \otimes_{\mathcal{O}} \mathcal{O}(E)$.

Basically, we know that $\text{div } f + \text{div } g = \text{div } fg$: we're doing a bit of handwaving here. But then we end up with the exact sequence

$$0 \to \mathcal{O}(D - p) \to \mathcal{O}(D) \to \varepsilon \to 0,$$

where ε is a one-dimensional module concentrated at p (it's also a residue field like $\kappa(p)$, but not canonically). When we write down the cohomology, we're working with a curve, which means that the dimensions h^2 and higher are zero for all of $\mathcal{O}(D - p), \mathcal{O}(D)$, and ε.

We know that ε is a one-dimensional module, so its global sections have dimension 1: $\boxed{h^1(\varepsilon) = 1}$. And its support has dimension 0 (a single point has dimension 0), so $H^1(\varepsilon)$ has dimension 0.

So now there are only two possibilities to account for the boxed 1: either we have

$$h^0(\mathcal{O}(D)) = h^0\mathcal{O}(D - p) + 1 \implies h^1(\mathcal{O}(D)) = h^1\mathcal{O}(D - p),$$

meaning the 1 is absorbed into the sequence at level h^0, or

$$h^0(\mathcal{O}(D)) = h^0\mathcal{O}(D - p) \implies h^1(\mathcal{O}(D)) = h^1\mathcal{O}(D - p) - 1$$

meaning the 1 is absorbed at level h^1.

Definition 203

Let $\mathcal{M}$ be a finite $\mathcal{O}$-module. The **Euler characteristic** $\chi(\mathcal{M})$ is defined via

$$\chi(\mathcal{M}) = \sum_q (-1)^q h^q(\mathcal{M}).$$

In the case of a curve, only the dimensions h^0 and h^1 can be nonzero, so

$$\chi(\mathcal{M}) = h^0(\mathcal{M}) - h^1(\mathcal{M}).$$

Note that this already requires the first and third results from above to be a consistent definition. But we can plug this in to both of the two cases above:

Corollary 204

For a smooth curve, we always have

$$\chi(\mathcal{O}(D)) = \chi(\mathcal{O}(D-p)) + 1.$$

Our next result shows that we can actually compute the Euler characteristic:

Definition 205

The **degree** of a divisor $D \sum_i r_i p_i$ is the sum of its coefficients:

$$\deg D = \sum_i r_i.$$

Theorem 206 (Riemann-Roch, version 1)

We have

$$\chi(\mathcal{O}(D)) = \deg D + \chi(\mathcal{O}).$$

So if there was no h^1, we would have already computed the dimension for our classical problem.

Proof. This theorem is clear for the zero divisor D. We can get to $\mathcal{O}(D)$ by adding or subtracting finitely many points.

But each time we go from D to $D - p$ or vice versa, and the above corollary shows us that we gain the appropriate 1 in each case. So induction of our previous corollary yields the result. □

This version of Riemann-Roch is good for h^0 (looking at the global sections), but not for h^1.

Corollary 207

Let f be a rational function. Then $\deg(\text{div}(f)) = 0$ (this means the number of zeros of f is equal to the number of poles of f), and this is also true if we replace "zeros of f" with "zeros of $f - c$" for some $c \in \mathbb{C}$.

Proof. Let $D = \text{div}(f)$, and consider the isomorphism $\mathcal{O}(D) \xrightarrow{f} 0$. Indeed, the sections of $\mathcal{O}(D)$ on U are the rational functions α such that $\text{div}(\alpha) + D \geq 0$ on U, and

$$\text{div}(f\alpha) = \text{div}(f) + \text{div}(\alpha) \geq 0$$

on all of U. Thus, $f\alpha$ must be a section of $\mathcal{O}$ on U. Going backwards, we know by Riemann-Roch that $\chi(\mathcal{O}(D)) = \deg D + \chi(\mathcal{O})$, but we also know that $\chi(\mathcal{O}(D)) = \chi(\mathcal{O})$, so $\deg D = 0$ because isomorphic modules have the same characteristic. And f has the same poles as $f - c$, so we also get the same zeros. □

Now notice that $H^0(\mathcal{O}) = \mathbb{C}$: the number of zeros is equal to the number of poles, but global sections have no poles, so there can't be any zeros either. Let $h^1(\mathcal{O})$ be the **arithmetic genus** p_a of $\mathcal{O}$: thus Riemann-Roch can be rewritten as

$$\chi(\mathcal{O}(D)) = \deg D + 1 - p_a.$$

And thus notice that

$$\lim_{n \to \infty} \dim H^0(\mathcal{O}(np)) = \infty,$$

because $\deg(np)$ goes to infinity, meaning $\chi(\mathcal{O}(D))$ goes to infinity, and $\chi = h_0 - h^1$. In other words, there exists a rational function f with a pole only at p.

Corollary 208

As a final example, we can show that any projective curve X is connected in the classical topology.

Proof. Suppose otherwise: say that p is an isolated point, and we can write $X = X_1 \cup X_2$ for clopen X_1, X_2. Pick a point $p \in X_1$, and choose a rational function f with a pole only at p (meaning there are no poles on X_2). f is analytic on X_2, which is a compact space, and a bounded analytic function on a compact manifold is some constant c (this is the maximum principle).

But then $f - c$ is zero on X_i, and it is a rational function on a curve, so it can only have finitely many zeros, which is a contradiction. □

29 April 27, 2020

Today, we're going to apply the Riemann-Roch theorem to curves of genus 0, 1, and 2. We'll start with the case where we have a smooth curve Y with function field K.

Proposition 209

Any point $(\alpha_0, \cdots, \alpha_n)$ of $\mathbb{P}^n$ with values in K defines a morphism $\pi : Y \to \mathbb{P}^n$.

Proof. For the morphism to be defined on all of Y, we must be able to pick an index i for each $q \in Y$ such that $\beta_j = \frac{\alpha_j}{\alpha_i}$ are all regular at q (we don't need to worry about being nonzero, since $\frac{\alpha_i}{\alpha_i} = 1$). Then we define the morphism via

$$\pi(q) = (\beta_0(q), \cdots, \beta_n(q)).$$

Because Y is a smooth curve, its local rings are valuation rings. Therefore, if we pick i so that the valuation at q, $v_q(\alpha_i)$, is minimal (highest order pole), dividing must yield something that is regular (because $v_q\left(\frac{\alpha_j}{\alpha_i}\right) = v_q(\alpha_j) - v_q(\alpha_i)$). □

Proposition 210

Let Y be a projective curve, and let $\pi : Y \to X$ be a nonconstant morphism. Then π is finite (meaning that the preimage $Y' = \pi^{-1}X'$ of any affine open set $X' = \operatorname{Spec} A$ is affine, and writing it as $\operatorname{Spec} B$ makes B a finite A-module).

Proof. The fibres of X are closed subsets of Y, but Y is a curve of dimension 1 and the whole curve doesn't map to a point. So our fibres are finite sets, and we can use Chevalley's Finiteness Theorem. □

Proposition 211

Let $\pi : Y \to X$ be a nonconstant finite morphism of curves, and let K and L be the function fields of X and Y, respectively. Then L is a finite field extension of K of some degree n, and the fibres of X have n elements for all but a finite number of points in X.

Proof sketch. We can assume that X, Y are their affine open sets, $X = \operatorname{Spec} A, Y = \operatorname{Spec} B$, and B is a finite A-module. Pick a primitive element of the field extension and look at the minimal polynomial of that element (which is of degree n): then the discriminant is nonzero at all but finitely many points in X. □

With that, we can move on to discussing Riemann-Roch.

- First of all, suppose Y has (arithmetic) genus 0, which means that the dimension $p_a = h^1\mathcal{O}_Y$ is 0. This arithmetic genus is also equal to the geometric genus g for projective smooth curves (which we'll show later on).

  Consider the exact sequence
  $$0 \to \mathcal{O} \to \mathcal{O}(p) \to \varepsilon \to 0,$$
  where ε is the one-dimensional module supported at p (isomorphic to the residue field module). Then the cohomology sequence has dimensions as follows because $h^0\mathcal{O}$ is just the constant functions:

	$\mathcal{O}$	$\mathcal{O}(p)$	ε
h^0	1	?	1
h^1	0	?	0
h^2	0	0	0.

  But we also have coboundary maps which help us figure out the rest of the table too:

	$\mathcal{O}$	$\mathcal{O}(p)$	ε
h^0	1	2	1
h^1	0	1	0
h^2	0	0	0.

  This means that $H^0\mathcal{O}(p)$ contains 1 as well as some nonconstant function α: by looking at the divisors, α has a pole of order 1 at p and no other poles. Consider the morphism
  $$\pi : Y \overset{(1,\alpha)}{\to} \mathbb{P}^1 = X,$$
  which evaluates at a point q to produce $f(q) = (1, \alpha(q))$ when $q \neq p$ and $(\alpha^{-1}(p), 1) = (0, 1)$ at $q = p$ (where we have a pole). Remember that Riemann-Roch shows that α has a simple pole and takes on every value exactly once (because the divisor has degree 0, meaning that the number of zeros of $f - c$ is the same as the number of poles of f), which means that the map $\pi : Y \to X$ is a **bijective** finite morphism. Since X and Y have equal function fields, but both are normal, this means that all curves of genus 0 are isomorphic to $\mathbb{P}^1$.

- In the case where we have genus 1, our table looks like the following:

	$\mathcal{O}$	$\mathcal{O}(p)$	ε
h^0	1	$*$	1
h^1	1	?	0

  There are two possibilities for $*$: either $(h^0(\mathcal{O}(p)), h^1(\mathcal{O}(p))) = (1, 0)$ or $(2, 1)$. But if $h^0(\mathcal{O}(p)) = 2$, then there is a nonconstant rational function α which means the curve is isomorphic to $\mathbb{P}^1$, a contradiction. Thus we must have the $(1, 0)$ case.

  Now we can consider the exact sequence
  $$0 \to \mathcal{O}((k-1)p) \to \mathcal{O}(kp) \to \varepsilon \to 0,$$

where we're working with a slightly different ε, and say that $k \geq 2$. We can inductively compute the dimension: putting in $(k-1,0)$ in the first column for $\mathcal{O}((k-1)p)$ means that we have a second column of $(k,0)$: in other words, $h^0\mathcal{O}(kp) = k$ for all $k \geq 2$. Then $H^0(\mathcal{O}(2p))$ contains 1, but it also contains a nonconstant function x such that x has a pole of exactly order 2 at p (it can't be order 1 by our earlier argument).

Next, we can consider $h^0(\mathcal{O}(3p))$: this also contains 1 and x, and it also contains an element y with a pole of order 3 (because it is in $H^0\mathcal{O}(3p)$ but not $H^0\mathcal{O}(2p))$I, which is independent from x because we have different order poles. We can consider the morphism

$$Y \overset{(1,x,y)}{\to} \mathbb{P}^2$$

sending a point q to $(1, x(q), y(q))$ whenever $q \neq p$. Because y has a bigger order pole, this means that at $q = p$, we send p to $(y^{-1}, xy^{-1}, 1) = (0,0,1)$. This is a finite morphism, so the image is some curve X in $\mathbb{P}^2$ (its dimension has to be 1).

Now consider some line $L : ax_0 + bx_1 + cx_2 = 0$ in $\mathbb{P}^2$, where a, b, c are generic. Then the points of Y that map to L are those $q \in Y$ such that $a + bx(q) + cy(q) = 0$, $a + bx + cy$ is a rational function on Y, and it has a pole of order 3 at p, so it takes on every value 3 times. Thus, there are three points q with $\pi(q) \in L \cap X$, meaning that X is a cubic curve: let's find its equation.

Notice that x^2 has a pole of order 4 at p, so it is in $\mathcal{O}(4p)$, and $1, x, y, x^2$ are still independent. Thus, the rational functions $\{1, x, y, x^2\}$ form a basis for $H^0\mathcal{O}(4p)$, and similarly $\{1, x, y, x^2, xy\}$ for $H^0\mathcal{O}(5p)$. Now y^2 and x^3 both have a pole of order 6 and we have a 6-dimensional space, so we get a linear dependence. This gives us a **plane cubic** X, so every curve of genus 1 is isomorphic to an elliptic curve.

- Finally, we have the following table for curves of genus 2:

	$\mathcal{O}$	$\mathcal{O}(p)$	ε
h^0	1	1	1
h^1	2	1	0

But now look at $\mathcal{O} \to \mathcal{O}(p) \to \mathcal{O}(2p) \to \varepsilon$:

	$\mathcal{O}(p)$	$\mathcal{O}(2p)$	ε
h^0	1	$*$	1
h^1	1	$*$	0

The middle column either looks like $(1,0)$ or $(2,1)$, but it turns out that both are impossible.

We'll move on soon to the second version of Riemann-Roch, which tells us more about what H^1 actually looks like: that will allow us to say more about what's going on in this last case!

30 April 29, 2020

As an interlude, we'll discuss the Birkhoff-Grothendieck theorem today:

Theorem 212 (Birkhoff-Grothendieck)

Let $X = \mathbb{P}^1$, and let $\mathcal{M}$ be a locally free $\mathcal{O}_X$-module of rank r. Then

$$\mathcal{M} \cong \bigoplus_{i=1}^{r} \mathcal{O}(n_i).$$

splits into a direct sum of (rank one) twisting modules.

This was proved by Grothendieck in this form, and it's equivalent to a version with matrices proved by Birkhoff: let $A_0 = \mathbb{C}[u]$, $A_1 = \mathbb{C}[v]$, where $u = \frac{x_1}{x_0}$, $v = \frac{x_0}{x_1}$ (these are the coordinate rings of the two standard affine open sets). Then A_{01}'s elements are the Laurent polynomials $\mathbb{C}[u, u^{-1}]$. Then we have the following:

Theorem 213

Suppose N is an invertible A_{01}-matrix. Then there exists an invertible A_0-matrix Q and an invertible A_1-matrix P such that $Q^{-1}NP$ is diagonal with diagonal entries that are integer powers of u.

When we look at a locally free module $\mathcal{M}$, we'll get a free module on U^0 and on U^1. Then N basically tells us how to glue those two together, and this idea above of changing basis tells us that we can do exactly that in a simple way.

We're going to prove Grothendieck's version of this here (we'll verify the Birkhoff matrix version on our homework for 2 × 2 matrices), but we'll need to review some terminology first.

Definition 214

Let A be a domain and M be an A-module. An element $m \in M$ is a **torsion element** if $am = 0$ for some $a \neq 0$ in A (that is, the annihilator of m is nonzero). The set of torsion elements in M forms a module called the **torsion submodule**. M is **torsion-free** if the submodule is zero, and it is a **torsion module** if all elements have torsion.

This works well with localization: if M is a torsion-free A-module and $s \neq 0$ is an element of A, then M_s is a torsion-free A_s-module. Thus, the definition of "torsion-free" extends to $\mathcal{O}$-modules as well.

We also want a few results about curves:

Lemma 215

Let X be a curve and $\mathcal{M}$ be an $\mathcal{O}_X$-module. If the torsion submodule is nonzero, then $\mathcal{M}$ has a global section: $H^0\mathcal{M} \neq 0$.

Proof. Replace $\mathcal{M}$ by its torsion submodule, and assume it is finitely generated (take a finitely generated submodule). This is okay, because we're just trying to show that H^0 is not zero.

Each element in $\mathcal{M}$ is killed by something, so its annihilator is a nonzero ideal, meaning that the annihilator module is nonzero. Thus, the support supp $\mathcal{M}$, which is a closed subset, is a finite set S (this is where we use the fact that X is a curve). So now we use the sheaf axiom: we can choose an affine open set U that contains S. Let $V = X - S$: then $X = U \cap V$ is a union of two open sets, and $\mathcal{M}(V) = 0$, so $\mathcal{M}(X) = \mathcal{M}(U)$. But U is affine, so $\mathcal{M}(U)$ is nonzero as long as the module isn't zero, and thus $\mathcal{M}(X)$ is nonzero. □

Proposition 216

Let X be a smooth curve and $\mathcal{M}$ be a torsion-free finite $\mathcal{O}$-module. Then $\mathcal{M}$ is locally free: that is, there exists an affine covering $\{U^i\}$ for X such that $\mathcal{M}(U^i)$ is isomorphic to $\mathcal{O}(U^i)^r$ for some r.

Proof. Since this is a local statement, we can assume $X = \operatorname{Spec} A$. Let p be a point in X: recall the local ring A_p is the set of elements $\alpha \in K$ in the fraction field of A that are regular at p, and we can get A_p by adjoining all elements a that don't evaluate to zero at p.

Because X is a smooth curve, we know that this local ring is a valuation ring. If $\mathcal{M}$ is a torsion-free $\mathcal{O}$-module, then this corresponds to a finite A-module M (because $\mathcal{O}$-modules and A-modules correspond). Then

$$\mathcal{M}_p = \mathcal{M} \otimes_A A_p$$

is the module obtained by localizing $\mathcal{M}$ at p. Since A_p is a valuation ring, it is a principal ideal domain, so every torsion free module is free: $\mathcal{M}_p$ is isomorphic to A_r^p by using a basis $(m_1, \cdots, m_r)$ of $\mathcal{M}_p$. Since there are finitely many of these, the m_i are all in some localizations M_s for a nonzero $s \in A$, and then we get a map $A_s^r \to M_s$, which we can embed in the exact sequence

$$0 \to K \to A_s^r \to \mathcal{M}_s \to C \to 0.$$

But we can localize this sequence: when we localize at a point p, we have that $K_p = C_p = 0$ (because of the isomorphism), so there exists some $s' \in A$ such that $K_{ss'} = C_{ss'} = 0$, and now we've shown the result we want (isomorphism for a localization). □

Definition 217

Let A be a Noetherian domain, and let M, N be A-modules. **$\mathrm{Hom}_A(M, N)$** is an A-module consisting of the module homomorphisms $M \to N$: addition and scalar multiplication of these homomorphisms are defined in the obvious way.

We'll write this as ${}_A(M, N)$. There are some relevant functorial properties here: if we have a module homomorphism $N_1 \to N_2$, we can map from ${}_A(M, N_1) \to_A (M, N_2)$. On the other hand, when we have a module homomorphism $\psi : M_1 \to M_2$, we have a **contravariant** map ${}_A(M_2, N) \to_A (M_1, N)$.

Fact 218

If $0 \to N_1 \to N_2 \to N_3$ is an exact sequence, then so is $0 \to_A (M_1, N_1) \to_A (M_2, N_2) \to_A (M_3, N_3)$. Meanwhile, if $M_1 \to M_2 \to M_3 \to 0$ is exact, then so is (again, contravariant) $0 \to_A (M_3, N) \to_A (M_2, N) \to_A (M_1, N)$.

Proposition 219 (Splitting)

If we have a set of maps $\mathcal{N} \xrightarrow{i} \mathcal{M} \xrightarrow{s} X$, then $\mathcal{M}$ is isomorphic to the direct sum $\mathcal{M} \cong \mathcal{N} \oplus K$, where $K = \ker s$.

(We just need to check that $\mathcal{N} \cap K$ is empty and that $\mathcal{N} + K = \mathcal{M}$.)

Corollary 220

Any functor F on $\mathcal{O}$-modules carries direct sums to direct sums.

One final fact to know about Hom is the following:

Lemma 221

Let M, N be A-modules for some Noetherian domain A. Let $s \in A$ be a nozero element: if M is a finite module, then

$${}_{A_s}(M_s, N_s) \cong ({}_A(M, N))_s.$$

We'll finish the proof next time: we'll want to refer to the **dual module**.

Definition 222

The **dual module** of an A-module M is

$$M^* =_A (M, A).$$

This is compatible with localization: for instance, if $\mathcal{M}$ is an O-module and $\mathcal{M}^* =_{\mathcal{O}} (\mathcal{M}, \mathcal{O})$ (because Hom is contravariant in the first variable), a map $\mathcal{M}_1 \to \mathcal{M}_2$ will give us a map $\mathcal{M}_2^* \to \mathcal{M}_1^*$. (And notice that we define the dual module in this order because $_A(A, M) \cong M$: the homomorphism is dependent on where 1 is sent.)

31 May 1, 2020

We'll put the Birkhoff-Grothendieck theorem aside for now and focus on a new topic: the second version of Riemann-Roch.

We're going to need some new terminology, and our first topic will be that of **branched covers** of curves. Suppose that $\pi : Y \to X$ is a nonconstant morphism of smooth projective curves: recall that Chevalley's Finiteness Theorem tells us that π is a finite morphism, so if $U = \operatorname{Spec} A$ is an affine open subset of X, then its preimage $\pi^{-1}U$ is also affine: writing it as $\operatorname{Spec} B$, we'll have an injective homomorphism $A \to B$, and B will be a torsion-free finite A-module, and therefore it will be locally free. This means that it makes sense to take the direct image $\pi_*\mathcal{O}_Y$, and this will be a finite $\mathcal{O}_X$-module (because the modules on the affine open sets are all finite) that is locally free, say of rank n. We'll just make this our definition:

Definition 223

A morphism $\pi : Y \to X$ of smooth projective curves is called a **branched cover**.

One important concept here is that of ramification: let $q \in Y$ have image $p = \pi(q)$. We have a local ring at each point of our smooth curve, so we can let x be the generator for our local ring of X at p (that is, it is a generator in some affine open set). Because p is the image of q, and the local rings are valuation rings, we can also evaluate x at q:

$$v_q(x) = e$$

for some $e \geq 0$, which we call the **ramification index at *q***. Remember that our fibres are finite, so each point $p \in X$ has some points $q_1, \cdots, q_k$ in the fibre over p.

Proposition 224

Let c_i be the ramification index at q_i. Then $\sum_{i=1}^k e_i = n$ is the rank of the $\mathcal{O}_X$-module $\pi_*\mathcal{O}_Y$, and the extended ideal can be written as

$$\mathfrak{m}_p\mathcal{O}_Y = \mathfrak{m}_{q_1}^{e_1} \cdots \mathfrak{m}_{q_k}^{e_k}.$$

In the case where we have affines $U = \operatorname{Spec} A$ and $\pi^{-1}U = \operatorname{Spec} B$, then $\mathfrak{m}_p$ is an ideal of A, and we have a map $A \to B$ such that

$$\mathfrak{m}_p B = \mathfrak{m}_{q_1}^{e_1} \cdots \mathfrak{m}_{q_k}^{e_k}.$$

We'll move on to another preliminary topic:

Definition 225

Let Y be a smooth curve. An $\mathcal{O}_Y$-module $\mathcal{L}$ is **invertible** if it is locally free of rank 1.

The reason for the name here is that

$$\mathcal{O}(D) \otimes_{\mathcal{O}} \mathcal{O}(E) \cong \mathcal{O}(D+E):$$

functions with poles D times functions with poles E are the functions with poles $D+E$. This means that

$$\mathcal{O}(D) \otimes_{\mathcal{O}} \mathcal{O}(-D) \cong \mathcal{O},$$

and since $\mathcal{O}$ can be thought of as a "unit" when doing an $\mathcal{O}$-tensor product, $\mathcal{O}(D)$ and $\mathcal{O}(-D)$ are inverses.

Proposition 226

Every invertible module is isomorphic to $\mathcal{O}_Y(D)$ for some divisor D.

D isn't necessarily unique here, so it's interesting to ask when $\mathcal{O}(D)$ and $\mathcal{O}(E)$ are isomorphic. Remember that the sections of this module on an open set U are the rational functions α on Y such that

$$\text{div}(\alpha) + D$$

(where the divisor div counts the zeros minus the poles) is an effective divisor (≥ 0) on U. (And we also need to include the element 0.) So suppose we have an isomorphism $\phi : \mathcal{O}(D) \to \mathcal{O}(E)$: if we look on an open set not containing any points in D, then we just want the sections of $\mathcal{O}$, the structure sheaf.

The idea is to note that $\mathcal{O}(D)$ and $\mathcal{O}(E)$ are both subsets of the function field of Y, which we denote K. "Lifting" the isomorphism to $K \to K$, we know that any map from K to itself is a multiplication by a rational function f. Thus, we need to find a function f such that

$$\text{div}(\alpha) + D \geq 0 \iff \text{div}(f\alpha) + E \geq 0,$$

but because the divisor is multiplicative, this means $\text{div}(f\alpha) = \text{div}(f) + \text{div}(\alpha)$, and then we can equivalently say that

$$D \geq 0 \iff \text{div}(f) + E \geq 0,$$

meaning that $\boxed{\text{div}(f) = D - E}$.

Definition 227

Two divisor D, E are **linearly equivalent** if $D - E = \text{div}(f)$ for some rational function f.

Corollary 228

$\mathcal{O}(D) \cong \mathcal{O}(E)$ if and only if D and E are linearly equivalent.

A final preliminary concept we're going to need is that of a **differential**:

Definition 229

Let A be an algebra (scalars $\mathbb{C}$), and let M be an A-module. A **derivation** is a map $\delta : A \to M$ such that $\delta(a+b) = \delta(a) + \delta(b)$, $\delta(ab) = a\delta(b) + b\delta(a)$, and $\delta c = 0$ for all scalars $c \in \mathbb{C}$.

(These are like the calculus rules for differentiation.)

Definition 230

A **module of differentials** Ω_A is an A-module generated by elements da, for each $a \in A$, with the relations $d(a+b) = da + db, d(ab) = a\,db + b\,da, dc = 0$ for $c \in \mathbb{C}$.

These relations cut down the module to a finite A-module. Since the relations look the same in the two above definitions, we can make the following observation:

Corollary 231

Derivations $\delta : A \to M$ correspond bijectively to module homomorphism $\phi : \Omega_A \to M$, such that

$$\delta = \phi \circ d$$

(where d maps A to ΩA).

Suppose, for example, that $R = \mathbb{C}[x_1, \cdots, x_n]$. Then Ω_R is a free R-module with basis $dx_1, \cdots, dx_n$, because any polynomial f has differential

$$df = \sum_i \frac{\partial f}{\partial x_i} dx_i,$$

and we can inductively continue this reasoning.

Note that Ω is compatible with localization: if $s \neq 0$ is an element of a domain A, then $\Omega_{A_s} \cong (\Omega_A)_s$. That means we can define the $\mathcal{O}_X$-module Ω_X by letting its sections on an affine open set $U = \operatorname{Spec} A$ be Ω_A.

One other fact: suppose that I is an ideal of a ring A, and $\overline{A} = A/I$. Consider the submodule

$$N = dI + I\Omega_A,$$

which tells us that $\overline{d\alpha} = 0$ for any $\alpha \in I$, as well as that $\overline{\alpha\gamma} = 0$ for any $\alpha \in I, \gamma \in \Omega_A$. Then we have that $\Omega_{\overline{A}} \cong \Omega_A / N$.

Proposition 232

Let Y be a smooth curve. Then Ω_Y is an invertible module.

Therefore, we can write $\Omega_Y \cong \mathcal{O}(K)$ for some divisor K, which we call the **canonical divisor** (up to linear equivalence).

With all of this, we're finally able to state our result:

Theorem 233 (Riemann-Roch, version 2)

Let Y be a smooth projective curve, let D be a divisor, and let K be a canonical divisor. Then

$$h^0\mathcal{O}(D) = h^1\mathcal{O}(K-D), \quad h^1(\mathcal{O}(D)) = h^0\mathcal{O}(K-D).$$

(These two equations are equivalent because we can plug in $K - D$ for D in the first equation.) It's not very clear why differentials come up, and the proof won't make that very clear either.

A corollary of this theorem will be the following result:

Theorem 234

For any smooth projective curve Y, we have $p_a = g$ (where $p_a = h^1\mathcal{O}_Y$ and g is the topological genus). In addition, the degree of the canonical divisor is $2g - 2 = 2p_a - 2$.

32 May 4, 2020

We'll start with a bit of review: recall that for a variety X, we can construct an $\mathcal{O}$-module of differentials Ω_X. If we write $X = \operatorname{Spec} A$, then Ω_A is generated by elements $\{da : a \in A\}$, such that $d(a+b) = da + db$, $d(ab) = adb + bda$, and $dc = 0$ for $a, b \in A$ and $c \in \mathbb{C}$. When Y is a msooth curve, Ω_Y is locally free of rank 1 (it is invertible), so we know by Birkhoff-Grothendieck that $\Omega_Y \cong \mathcal{O}(K)$ for some **canonical divisor** K. (Note that we can have different canonical divisors if they are linearly equivalent – that is, their difference is the divisor of a rational function.)

Then the second version of Riemann-Roch tells us that if Y is a smooth projective curve and D is a divisor, then

$$h^0(\mathcal{O}(D)) = h^1(\mathcal{O}(K - D))$$

(and therefore the same result holds when we swap the roles of D and $K - D$, since we can plug in $D \to K - D$).

Note that when we apply this to the divisor $D = 0$, $h^0\mathcal{O}(K) = h^1\mathcal{O}$ is the genus g of the variety, and $h^1\mathcal{O}(K) = h^0\mathcal{O} = 1$. Thus, the Euler characteristic $h^0 - h^1$ here is $\chi(\mathcal{O}(K)) = g - 1$.

Corollary 235

The degree of the canonical divisor D is $2g - 2$.

Proof. We know that

$$\chi(\mathcal{O}(K)) = \deg K + (1 - g)$$

by Riemann-Roch version 1, and then the left hand side is $g - 1$ by the argument above. □

Example 236

Let's consider a curve of genus $g = 2$.

Then the degree of the canonical divisor, $\deg K$, is $2g - 2 = 2$, and the Euler characteristic $\chi\mathcal{O}(K) = 1$. This means that $h^0\mathcal{O}(K) > 0$, so there exists some nonzero global section, which means that there exists a rational function α such that

$$\operatorname{div}(\alpha) + K \geq 0.$$

But now the left hand side is $\operatorname{div} K'$ for some canonical divisor K' (K and K' are linearly inequivalent), so we can **assume that** K' **is effective** and of the form $p_1 + p_2$, where it's possible that $p_1 = p_2$.

Riemnann-Roch now tells us that

$$h^0(\mathcal{O}(D)) = h^1(\mathcal{O}) = g = 2,$$

so H^0 has some basis $(1, \alpha)$ (we can always choose 1 to be a global section), such that α has poles at p_1 and p_2. This means α has order 2, so takes on every value 2 times (with multiplicity).

But now $(1, \alpha)$ is a point of $\mathbb{P}^1$ with values in the function field F of Y, and any such point defines a morphism $Y \to \mathbb{P}^1$ via $\pi(q) = (1, \alpha(q))$ whenever $q \neq p_1, p_2$ and $(\frac{1}{\alpha(q)}, 1) = (0, 1)$ when $q = p_1, p_2$. We already mentioned that π is a double covering of $\mathbb{P}^1$ above.

Definition 237

A smooth projective curve Y is **hyperelliptic** if it can be represented as a double cover of $\mathbb{P}^1$.

Every curve of genus 1 can be represented as such a covering, so this may explain the "elliptic" part of the definition. One way to create such a curve is to consider $y^2 = f(x_0, x_1)$ in the weighted projective plane.

Corollary 238

Every smooth projective curve Y of genus 2 is hyperelliptic.

Example 239

Now let's consider $g = 3$.

We can again assume that K is effective, so it is the sum of four points, and the global sections form a space of dimension $h^0\mathcal{O}(K) = g = 3$: suppose we have a basis $(1, \alpha, \beta)$. Then α, β must have poles in the set $\{p_1, p_2, p_3, p_4\}$.

We claim that all points must be represented. Suppose otherwise, so that both α and β are sections of $\mathcal{O}(p_1 + p_2 + p_3) = \mathcal{O}(K - p)$, where $p = p_4$. Then

$$h^0\mathcal{O}(p_1 + p_2 + p_3) = h^0\mathcal{O}(K - p) = h^1\mathcal{O}(p)$$

by Riemann-Roh version 2. But now we can estimate the cohomology of $\mathcal{O}_p$ via the exact sequence

$$0 \to \mathcal{O} \to \mathcal{O}(p) \to \varepsilon \to 0.$$

We know that $\mathcal{O}(p)$ does not have a nonconstant global section, so the table of cohomology dimensions looks like

	$\mathcal{O}$	$\mathcal{O}(p)$	ε
h^0	1	1	1
h^1	3	2	0

But then $h^1\mathcal{O}(p)$ is not 3, which is a contradiction. So we can assume that all points are represented, and now for any point $q \neq p_1, p_2, p_3, p_4$, we can construct the point with values in K $(1, \alpha, \beta)$. The only points where α and β may not be regular are the p_is, so there are only the four points p_i that are on the line $X_0 = 0$. So the image of Y is a plane curve, and we have a finite morphism $YtoX$ such that the degree of X is the number of points in $L \cap X$.

Since 4 points map to L, this means we either have X of degree 1 and a covering Y of degree 4, X of degree 2 and a covering of degree 2, or X of degree 4 and a generically injective covering. But the first case can't happen because $(1, \alpha, \beta)$ are linearly independent.

The other two cases can be be described: if X is degree 2, then X is a conic isomorphic to $\mathbb{P}^1$, and Y is a double covering and therefore hyperelliptic. And finally, if X is degree 4, then Y is the **normalization** of X, since the genus of X has dimension $\binom{4-1}{2} = 3$.

Corollary 240

A smooth projective curve of genus 3 is either hyperelliptic or isomorphic to a plane curve of degree 4.

We'll finish by verifying Riemann-Roch in the case where $Y = \mathbb{P}^1$:

Proof. If we apply Birkhoff-Grothendieck to invertible modules, we see that every invertible $\mathcal{O}$-module is isomorphic to $\mathcal{O}(n)$ for some n, which we wish to identify.

Identify Ω_X with $\mathcal{O}(K)$, where $X = U^0 \cup U^1$. Identify $U^0 = \operatorname{Spec} \mathbb{C}[u]$ and $U^1 = \operatorname{Spec} \mathbb{C}[v]$. We know that Ω_{U^0} is free with basis du, and we know that $du = dv^{-1} = -v^{-2}dv$, so du has a pole of order 2 at the point at infinity, meaning du is a global section of $\Omega_X(2p_\infty)$. But this section is zero nowhere, so $\Omega_X(2p_\infty)$ is free and isomorphic to $\mathcal{O}_X$, meaning that

$$\mathcal{O}_X(-2p_\infty) \cong \mathcal{O}(-2).$$

So now we can verify Riemann-Roch by noting that $\mathcal{O}(D)$ is isomorphic to $\mathcal{O}(n)$ and that $\mathcal{O}(K - D)$ is isomorphic to $\mathcal{O}(-n-2)$. We're supposed to have $h^0\mathcal{O}(n) = h^1\mathcal{O}(--2)$ and $h^1\mathcal{O}(n) = h^0\mathcal{O}((n-2)$.

When $n \geq 0$: then $h^0\mathcal{O}(n) = n + 1$ and $h^1\mathcal{O}(n) = 0$ (because we know the cohomology of the twisting modules). Meanwhile, when $r > 0$, $h^0\mathcal{O}(-r) = 0$ and $h^1\mathcal{O}(-r) = -r - 1$. Plugging in $r = n + 2$, we find that

$$h^1\mathcal{O}(-n-2) = n + 1, \quad h^0\mathcal{O}(-n-2) = 0,$$

and indeed things work out. □

33 May 6, 2020

Today, we'll want to generalize Riemann-Roch from $\mathbb{P}^1$ to curves in general. Proving it for $\mathbb{P}^1$ is just the computation that we did in class, and in general we will want to make Y into a branched covering of $X = \mathbb{P}^1$. Then $\pi : Y \to X$ is a finite morphism of some degree n, meaning that

$$H^q(Y, \mathcal{O}_Y(D)) = H^q(X, \pi_*\mathcal{O}_Y(D)),$$

so we can get an analogous statement for $\mathcal{O}_Y(D)$ by taking its direct image and applying Riemann-Roch. But there are two main problems with this: first of all, $\pi_*\mathcal{O}_Y(D)$ is not invertible (unlike $\mathcal{O}_Y(D)$), though it is locally free of rank n. This isn't too much of a problem, since we can use the Birkhoff-Grothendieck theorem to write the locally free module as a direct sum of $\mathcal{O}(r_i)$s and then apply Riemann-Roch there. But we have a bigger issue in that we need to relate $\pi_*\mathcal{O}_Y(K - D)$ to $\pi_*\mathcal{O}_Y(D)$, so we will need to rewrite some things here.

We'll write $\mathcal{O}_Y(K) = \Omega_Y$ as the module of differentials, and $\mathcal{O}_Y(D) = \mathcal{L}$ as an invertible $\mathcal{O}_Y$-module. Then the dual module

$$\mathcal{L}^* =_{\mathcal{O}_Y} (\mathcal{L}, \mathcal{O}_Y) \cong \mathcal{O}_Y(-D).$$

Lemma 241

Let $\mathcal{M}$ and $\mathcal{N}$ be $\mathcal{O}_Y$-modules, and let D be a divisor. Let ${}_{\mathcal{O}_Y}(\mathcal{M},\mathcal{N})$ be the $\mathcal{O}$-module with sections equal to the homomorphisms $\mathcal{M}(U) \to \mathcal{N}(U)$. Then

$${}_{\mathcal{O}_Y}(\mathcal{M},\mathcal{N}) ={}_{\mathcal{O}_Y} (\mathcal{M}(D),\mathcal{N}(D))$$

(where $\mathcal{M}(D) = \mathcal{M} \otimes_{\mathcal{O}} \mathcal{O}(D)$), and $\mathcal{O}(K - D)$ is isomorphic to ${}_{\mathcal{O}_Y}(\mathcal{L},\Omega_Y)$.

Proof. For the first part, we are given homomorphisms $\mathcal{M} \to \mathcal{N}$, tensoring with $\mathcal{O}(D)$ gives us a homomorphism $\mathcal{M}(D) \to \mathcal{N}(D)$. This is an invertible operation because we can tensor with $\mathcal{O}(-D)$ to get back to where we started.

For the second part, remember that ${}_{\mathcal{O}}(\mathcal{O},\mathcal{M}) \cong \mathcal{M}$ (because sections m of $\mathcal{M}$ correspond to the homomorphisms that are multiplication by m).

$${}_{\mathcal{O}_Y}(\mathcal{L},\Omega_Y) ={}_{\mathcal{O}_Y} (\mathcal{O}_Y(D),\mathcal{O}_Y(K)),$$

and tensoring both sides by $\mathcal{O}(-D)$ yields

$$\cong \mathcal{O}_Y(\mathcal{O}_Y,\mathcal{O}_Y(K - D)).$$

□

Definition 242

Let $\mathcal{M}$ be a locally free $\mathcal{O}_Y$-module. Then the **Serre dual** of $\mathcal{M}$, denoted $\mathcal{M}^{\#}$, is the set of maps from $\mathcal{M}$ to the differentials:

$$\mathcal{M}^{\#} ={}_{\mathcal{O}_Y} (\mathcal{M},\Omega_Y).$$

For example,

$$\mathcal{O}_Y(D)^{\#} ={}_{\mathcal{O}_Y} (\mathcal{O}(D),\mathcal{O}(K)) \cong_{\mathcal{O}_Y} (\mathcal{O},\mathcal{O}(K - D)) \cong \mathcal{O}_Y(K - D).$$

This tells us how to state Riemann-Roch version 2 in another way:

Theorem 243

Let $\mathcal{M}$ be a locally free $\mathcal{O}_Y$-module. Then

$$h^0\mathcal{M} = h^1\mathcal{M}^{\#}, \quad h^1\mathcal{M} = h^0\mathcal{M}^{\#}.$$

In order to apply the $\mathbb{P}^1$-version of Riemann-Roch 2, we need to show that

$$\pi_*(\mathcal{M}^{\#}) \cong (\pi_*\mathcal{M})^{\#},$$

where $\mathcal{M}^{\#}$ is the Serre dual on Y, and the right side is the Serre dual on X. We'll drop the π_* symbol from here: we'll denote the direct image of an $\mathcal{O}_Y$-module $\mathcal{M}$ on X by $\mathcal{M}$ as well. In other words, if $U = \operatorname{Spec} A$ is an affine open in X, and $V = \pi^{-1}U = \operatorname{Spec} B$ is affine, then the B-module $M = \mathcal{M}(V)$ can be viewed as an A-module because we have restriction of scalars through the map $A \to B$, and we'll say that the sections of the direct image on U are also denoted by M.

To make progress, we'll need to talk about the **trace** of a differential. If we have a map $\mathcal{O}_X \to \mathcal{O}_Y$ (which is the direct image), then the **trace** is the map $\mathcal{O}_Y \to \mathcal{O}_X$ such that the composition is multiplication by the degree of the covering. We'll let the trace of β be denoted as $\operatorname{tr}(\beta)$.

To understand the trace for differentials, we'll do a bit of computation: let $p \in X$ be a point, and let x be a local generator for the maximal ideal $\mathfrak{m}_p$. Then we pick $q \in Y$ such that $\pi(q) = p$, and let y be the local generator for $\mathfrak{m}_q$, we can let the ramification index be e (meaning that $x = uy^e$, and u is a local unit). Then

$$dx = y^e du + ey^{e-1} u dy,$$

but dy is a local generator for the module of differentials Ω_Y, and u is a regular function on Y at q. Therefore, du can also be written as wdy for some regular function w at q. Therefore,

$$dx = (wy^e + ey^{e-1}u)dy = vy^{e-1}dy,$$

where $v = wy + eu$ is a local unit (invertible function locally at q), meaning that

$$dy = \frac{1}{vy^{e-1}} dx$$

(so we have a pole of order $e - 1$ when we write things in terms of dy).

So now we can look on a fibre $q_1, \cdots, q_k$ of $p \in X$. If β is a differential on Y that is regular at all q_i, then we can write

$$\beta = bdx$$

(using the above local computation), where b has poles of order at most $e_i - 1$ at each point p_i.

Definition 244

The **trace** is a map τ of modules of differentials $\Omega_Y \to \Omega_X$, such that

$$\tau(\beta) = (\text{tr } b)dx.$$

Proposition 245 (Main Lemma)

If β is a differential on Y which is regular at $q_1, \cdots, q_k$ (the fibre over p), then $\tau(\beta)$ is a regular differential on X at p.

Proof. We know that $\beta = bdx$, where b has poles of order at most $e_i - 1$ at q_i. But we know that x has zeros of order e_i at q_i, so xb is regular at all q_i and evaluates to zero. Then

$$\tau(x\beta) = \text{tr}(xb)dx,$$

and the trace $\text{tr}(xb)$ is regular at p: in fact, we have a general formula. If f is a regular function on Y, then

$$[\text{tr}(f)](p) = \sum_i e_i f(q_i).$$

So if xb is zero at q_i, then $\text{tr}(xb) = 0$ at p, meaning that $\frac{1}{x}\text{tr}(xb)$ is regular at p as well. And because the trace is $\mathcal{O}_X$-linear, this last expression is equal to the trace of b, meaning

$$\tau(\beta) = \text{tr}(b)dx$$

is indeed regular at p, as desired. □

This trace defines a map $\Omega_Y \to \Omega_X$: note that

$$\Omega_Y \cong_{\mathcal{O}_Y} (\mathcal{O}_Y, \Omega_Y),$$

where the correspondence sends β to the map $b \to b\beta$. Then we can compose this map Ω_Y with τ to get to

$${}_{\mathcal{O}_X}(\mathcal{O}_Y, \Omega_X)$$

(because the first map is $\mathcal{O}_Y$-linear, but the second map is $\mathcal{O}_X$-linear).

Lemma 246

The map

$$\pi_*({}_{\mathcal{O}_Y}(\mathcal{O}_Y, \Omega_Y) \overset{\tau\circ}{\to} \mathcal{O}_X(\pi_*\mathcal{O}_Y, \Omega_X)$$

is bijective.

So differentials on Y are the same as maps from $\mathcal{O}_Y$ to Ω_X, which is a strange fact. As an exercise, we can look at the case where Y is defined by $y^3 = x$ and X is the affine line. Unfortunately, the proof is a bit long, and we can read it on our own.

So now if $\mathcal{M}$ is a locally-free $\mathcal{O}_Y$-module, then

$${}_{\mathcal{O}_Y}(\mathcal{M}, \Omega_Y) \overset{\tau\circ}{\to}_{\mathcal{O}_X} (\mathcal{M}, \Omega_X),$$

because we can look at an open set and treat $\mathcal{M}$ as a direct sum of $\mathcal{O}_Y$s, and then we can apply the above lemma. So now the left side is the Serre dual $\mathcal{M}^\#$ computed on Y, while the right side is $\mathcal{M}^\#$ computed on X, and we've shown Riemann-Roch version 2.

34 May 8, 2020

In these last two classes, we'll focus on applications of the Riemann-Roch theorem.

Proposition 247

The arithmetic genus p_a and topological genus g of a smooth curve are the same.

Proof. Recall that Riemann-Roch version 2 tells us that

$$h^0\mathcal{O}(D) = h^1\mathcal{O}(K - D), \quad h^1\mathcal{O}(D) = h^0\mathcal{O}(K - D).$$

Applying this with $D = 0$ yields

$$1 = h^0\mathcal{O} = h^1\mathcal{O}(K)$$

and

$$p_a = h^1\mathcal{O} = h^0\mathcal{O}(K).$$

But version 1 also tells us that

$$p_a - 1 = \chi(\mathcal{O}(K)) = \deg K + \chi(\mathcal{O}) = \deg K + 1 - p_a,$$

so the degree of the canonical divisor must be $2p_a - 2$.

We'll also compute this in another way. Suppose we have a branched covering $\pi : Y \to X = \mathbb{P}^1$, and Y has branch points q_i with ramifications e_i. Then looking at the differential dx on Y, if we let y_i be a local generator for $\mathfrak{m}_{q_i}$, then

$$x = uy_i^{e_i},$$

meaning that dx has a zero of order $e_i - 1$ at the point q_i. We'll asume that we're looking at the points over ∞ on $\mathbb{P}^1$ (this is just a change of coordinates): since dx has a pole of order 2 at the point ∞ on X, it must have a pole of order 2 at all of the point $p_1, \dots, p_n$ in the fibre over ∞.

So now let K be the divisor of dx on Y: this will be the sum of the zeros minus the sum of the poles, which takes the form

$$\sum_i (e_i - 1)q_i - 2\sum_{i=1}^n p_i.$$

Thus, the degree of K is $\sum(e_i - 1) - 2n$, which is also $2p_a - 2$ by our above calculation. But we are trying to work with the topological Euler characteristic of Y as well: we know that

$$e(Y) = ne(X) - (\text{number of times the different sheets come together}),$$

and this can be written in terms of the ramification indices: since e_i sheets come together at a point q_i, we lose $(e_i - 1)$. Thus,

$$e(Y) = n \cdot 2 - \sum_i (e_i - 1) = 2n - \sum_i (e_i - 1) = 2 - 2p_a$$

by consulting the calculation above. But in a manifold of degree 2, the Euler characteristic is also $2 - 2g$, so we must have $g = p_a$, as desired (and this tells us that $\deg K = 2g - 2$). □

We'll now move on to a different topic:

Definition 248

Let D be a divisor, and let $h^0\mathcal{O}(D) > 0$ (so we have some global sections). A point $p \in Y$ is a **base point** of $\mathcal{O}(D)$ if

$$h^0\mathcal{O}(D - p) = h^)\mathcal{O}(D).$$

In other words, whenever we have a function f such that $\text{div}(f) + D \geq 0$, we also have $\text{div}(f) + D - p \geq 0$. We can understand this using the usual exact sequence

$$0 \to \mathcal{O}(D - p) \to \mathcal{O}(D) \to \varepsilon \to 0,$$

which we use when proving Riemann-Roch version 1. When p is a base point, the h^0 dimensions are the same, which means that the h^1 values for the two modules differ by 1. Otherwise, h^0 will differ by 1.

In the case where D is effective, we can rephrase by saying that whenever p is not in the support supp D, if a function f has poles at most D, then f has poles at most $D - p$.

Example 249

The module $\mathcal{O}(K)$ has no base points.

We can see this by computing $h^0\mathcal{O}(K - p)$: by Riemann-Roch,

$$h^0\mathcal{O}(K - p) = h^1\mathcal{O}(p), \quad h^1\mathcal{O}(K - p) = h^0\mathcal{O}(p).$$

If Y is not the projective line, then $h^0\mathcal{O}(p)$ must be 1 – there is no function with just one simple pole. So

$$h^0\mathcal{O}(p) = h^0\mathcal{O} \implies h^0\mathcal{O}(K-p) = h^0\mathcal{O}(K) - 1,$$

which means that our arbitrary point p is not a base point.

Example 250

On the other hand, p is always a base point of $\mathcal{O}(K+p)$: every function with poles $K+p$ doesn't actually have the extra pole at p.

This is because

$$h^1\mathcal{O}(K+p) = h^0\mathcal{O}(-p) = 0 = h^0\mathcal{O} - 1,$$

and if the h^0 dimensions differ, the h^1 dimensions must be the same.

So now suppose that $D > 0$, and we use a basis of $H^0\mathcal{O}(D)$ (of regular functions) to apply the morphism

$$\pi : Y \to \mathbb{P}^r$$

via $(\alpha_0, \cdots, \alpha_r)$, where $h^)\mathcal{O}(D) = r+1$. Then remember that the **degree** of our morphism π is the number of points in $\pi^{-1}H$, where H is some generic hyperplane of $\mathbb{P}^r$.

Lemma 251

If D is effective and $\mathcal{O}(D)$ has no base points, then $\deg \pi = \deg D$.

Proof. Since H is generic, we'll represent it in the form

$$\sum c_i x_i = 0.$$

Then the preimage of H is the set of points

$$\{q \in Y : \sum c_i \alpha_i(q) = 0\}.$$

Let β be the regular function

$$\beta = \sum c_i \alpha_i :$$

since the α_is form a basis, this is also a global section in $H^0\mathcal{O}(D)$. The set of zeros of β is exactly the set $\pi^{-1}H$ that we want, but β will have all poles in the divisor D, since there are no base points in $\mathcal{O}(D)$. Thus, the number of poles is $\deg D$, and this is equal to the number of zeros of β as well. □

In the next topic that we cover, we'll let $g \geq 2$. Then we can use a basis of $H^0\mathcal{O}(K)$ (of size g) to map $\pi : Y \to \mathbb{P}^{g-1}$: this is known as the **canonical map**.

Theorem 252

The canonical map π is either an embedding of Y into $\mathbb{P}^{g-1}$, or X, the image of Y, is isomorphic to $\mathbb{P}^1$ and $Y \to X$ is a double cover (meaning that Y is hyperelliptic).

Proof. Suppose the map $\pi : Y \to X$ is not injective, so we have two points q_1, q_2 such that $\pi(q_1) = \pi(q_2)$. Consider $\mathcal{O}(K), \mathcal{O}(K-q_1)$, and $\mathcal{O}(K-q_1-q_2)$. We know that K is only chosen up to linear equivalence, so we can assume K is effective, and we can assume q_1, q_2 are not in the support of K.

Then $\mathcal{O}(K)$ has no base points, meaning

$$h^0\mathcal{O}(K - q_1) = h^0\mathcal{O}(K) - 1.$$

In other words, sppose our basis for $H^0\mathcal{O}(K)$ is $(\alpha_0, \cdots, \alpha_{g-1})$. Since this defines a morphism, these functions don't all vanish at q_1, but the dimension of functions that do vanish is only 1 less than the total dimension. So let's assume that $\alpha_0(q_1)$ does not vanish, but all the other α_is do: then $\pi(q_1) = (1, 0, \cdots, 0) = \pi(q_2)$ (by the initial choice of q_1, q_2), which means $\alpha_1, \cdots, \alpha_{g-1}$ alsovanish at q_2 (and $\alpha_0(q_2) \neq 0$ as well). So now

$$h^0\mathcal{O}(K - q_1 - q_2) = h^0\mathcal{O}(K - q_1),$$

and also

$$h^0\mathcal{O}(K - q_1 - q_2) = h^0\mathcal{O}(K - q_2) = g - 1,$$

because q_2 is a base point of $h^0\mathcal{O}(K - q_1)$. So now Riemann-Roch version 2 restates this as

$$h^1\mathcal{O}(q_1 + q_2) = h^1\mathcal{O}(q_1) = g - 1,$$

which means that

$$\chi(\mathcal{O}(q_1 + q_2)) = 2 + (1 - g) = 3 - g.$$

But $h^1(\mathcal{O}(q_1 + q_2)) = g - 1$, so $h^0 = 2$, which means that there is a two-dimensional space of functions with the poles q_1, q_2. Therefore, the map $Y \to \mathbb{P}^1$ where we use a two-dimensional basis of $h^0(\mathcal{O}(q_1 + q_2))$ has degree 2, which means Y is hyperelliptic.

Throughout all of this, we've assumed $Y \to X$ is not injective, so the other case is that $Y \to X$ is bijective. It's still possible that X is a cusp curve, but we'll skip over that – the way to deal with this case is to consider $2q$ instead of $q_1 + q_2$ and look at this reasoning. □

Theorem 253

Suppose Y is a curve of genus $g \geq 2$. Then there is at most one way to represent Y as a double cover of $\mathbb{P}^1$.

We know that if Y is hyperelliptic, there is a way to do this, but this theorem tells us that the representation is unique. We don't have time to prove this, but a sketch of the main ideas is that when we have a double cover $\pi : Y \to X = \mathbb{P}^1$, we can count the number of branch points by noting that

$$2 - 2g = 2e(X) - (\text{branch points}) = 4 - B,$$

meaning that $B = 2g + 2$. So if we look at the divisor of dx on Y, we have zeros of order $2 - 1 = 1$ at each of the branch points $q_1, \cdots, q_B$, and we have poles of order 2 at p_1, p_2 (the points in the fibre over ∞). Let K be the divisor of dx on Y, which can be written as

$$\sum_{i=1}^{B} q_i - 2(p_1 + p_2).$$

35 May 11, 2020

We'll discuss the **canonical embedding** of a curve today. Say that Y is a smooth projective curve: remember that Y is **hyperelliptic** if there exists a degree 2 map $Y \to X = \mathbb{P}^1$, meaning that Y is a double cover of X. As we discussed last time, there is at most one way to represent Y as a double cover of X when $g \ge 2$.

It's easy to construct hyperelliptic curves of any genus g: start with a homogeneous polynomial of degree $2d$ in x_0, x_1, and form the double line

$$y^2 = f(x),$$

where y is a weight d variable. Then the map $Y \to X$ ramifies at the zeros of f, meaning that the Euler characteristic

$$e(Y) = 2e(X) - 2d = 4 - 2d.$$

(because each of the $2d$ zeros is a branch point) But $C(Y) = 2 - 2g$, so $g = d - 1$

Another relevant fact is that whenever Y has genus ≥ 2 and is **not** hyperelliptic, $\mathcal{O}(K)$ has no base points, so its global sections define a map $\pi : Y \to \mathbb{P}^{g-1}$. This map is then an embedding of Y as a closed subvariety of projective space of degree g, and the degree of this embedding is $\deg K = 2g - 2$.

When the genus is 2, Y is always hyperelliptic, because $\mathcal{O}(K)$ will have degree 2, meaning the map $Y \to \mathbb{P}^1$ has degree $k = 2g - 2 = 2$, which is a double cover. We will proceed by describing this embedding in a few cases: $g = 3, 4, 5$.

Example 254

When $g = 3$, we are embedding Y into $\mathbb{P}^2$, so Y is a plane curve of degree $\deg K = 2g - 2 = 4$.

Indeed, this works out because the arithmetic and geometric genus are equal to

$$g = p_a = \binom{4-1}{2} = 3.$$

But there's another approach we can take here: take $R = \mathbb{C}[x_0, \cdots, x_n]$ and impose a grading

$$R = \bigoplus R_d,$$

where R_d consists of the homogeneous polynomials of degree d. The dimension of R_d is $\binom{d+n}{n}$, so it is $\binom{d+2}{2}$ in the case where we have three variables.

So now let $\alpha = (\alpha_0, \alpha_1, \alpha_2)$ be a basis for the global sections $H^0\mathcal{O}(K)$. Let A_1 be the span of these elements, let A_2 be the span of the degree-2 terms in α_is, and so on: this means we have a grading on A as well of the form

$$A = \bigoplus A_d.$$

Now $A_1 = H^0\mathcal{O}(K)$, and $A_2 \subset H^0\mathcal{O}(2K)$ and so on (because adding K gives an effective divisor for each α_i, so adding $2K$ will do so if we have $\alpha_i\alpha_j$). Now Riemann-Roch tells us that whenever $n \ge 2$,

$$h^0\mathcal{O}(nK) = \chi\mathcal{O}(nK) = \deg(nK) + 1 - g = 4n + 1 - g,$$

because $h^1\mathcal{O}(nK) = 0$ by Riemann-Roch version 2. The dimensions of R_d are $\binom{d+2}{2} = 1, 3, 6, 10, 15, 21, \cdots$ for $d = 0, 1, 2, 3, 4, 5$, and the dimensions of $h^0\mathcal{O}(dK)$ are $1, 3, 6, 10, 14, 18$ (by plugging in the above formula). Thus, there exists some nonzero homogeneous polynomial $f(x)$ of degree 4 such that $f(\alpha) = 0$, and there exist three homogeneous polynomials h_1, h_2, h_3 of degree 5, but that's not new information because they're just $x_0 f, x_1 f, x_2 f$. But now Y is contained in the zero locus of this degree 4 polynomial f, so indeed we've seen again that every non-hyperelliptic curve of genus 3 is a plane curve of degree 4.

Example 255

When $g = 4$, we are embedding Y into $\mathbb{P}^3$, and the degree of the embedding is $4 \cdot 2 - 2 = 6$.

We'll count dimensions again: this time $\dim R_d = \binom{d+3}{3}$, and $\dim A_d = \deg(dK) + 1 - g = 6d + 1 - g$ for all $d \geq 2$. Then $\dim R_d = 1, 4, 10, 20$ and $\dim A_d = 1, 4, 9, 15$ for $d = 0, 1, 2, 3$, which means that there exists a homogeneous quadratic f with $f(\alpha) = 0$. Then there are five cubics including $x_0 f, x_1 f, x_2 f, x_3 f$: call the last one g. Then Y is contained within the locus $Z = \{f = g = 0\}$, and now a version of Bezout's theorem (in higher dimensions) says that the degree of Z is $2 \cdot 3 = 6$. So every non-hyperelliptic curve of genus 4 is a **complete intersection** (it's defined via $f = g = 0$ in $\mathbb{P}^3$).

Example 256

When $g = 5$, we embed Y into $\mathbb{P}^4$ with $\deg K = 2g - 2 = 8$.

Then $\dim R_d = \binom{d+4}{4}$ is $1, 5, 15$ for $d = 0, 1, 2$, while $h^0\mathcal{O}(2K) = \deg(2K) + (1 - g) = 2(2g - 2) + (1 - g) = 12$. Thus, there are linearly independent three quadrics Q_1, Q_2, Q_3 that all contain Y.

Now Bezout's theorem tells us that the degree of $Q_1 \cap Q_2 \cap Q_3$ is $2 \cdot 2 \cdot 2 = 8$, which will also be the degree of Y if the dimension of the intersection $Q_1 \cap Q_2 \cap Q_3$ is 1. So then Y will be the intersection of 3 quadrics.

There's also another case above where the dimension of $Q_1 \cap Q_2 \cap Q_3$ is not of dimension but of dimension 2, and this case does happen. If we can represent Y as a triple cover of $\mathbb{P}^1$, then $Q_1 \cap Q_2 \cap Q_3$ will fall into his case, and Y is called **trigonal**.

Remark 257. *And similarly, for $g = 6$, $\deg Y = 10$ doesn't factor enough for Y to be a complete intersection.*