

Affine algebraic geometry

Closed algebraic sets in affine space

We assume that the reader has basic familiarity with the theory of dimension (i.e., Krull dimension) for Noetherian rings, and we also assume the Hilbert basis theorem, which asserts that a finitely generated algebra over a Noetherian ring is Noetherian. In particular, a finitely generated algebra over a field K is Noetherian. We also assume familiarity with Hilbert's Nullstellensatz.

Throughout, K denotes a field. We eventually focus on the case where K is algebraically closed, but at the outset, and, occasionally, elsewhere, we relax this assumption. The reader is welcome to think primarily of the case where K is the field \mathbb{C} of complex numbers, although for the purpose of drawing pictures, it is easier to think about the case where $K = \mathbb{R}$.

Let K be a field. A polynomial in $K[x_1, \dots, x_n]$ may be thought of as a function from $K^n \rightarrow K$. Given a finite set f_1, \dots, f_m of polynomials in $K[x_1, \dots, x_n]$, the set of points where they vanish simultaneously is denoted $\mathcal{V}(f_1, \dots, f_m)$. Thus

$$V(f_1, \dots, f_m) = \{(a_1, \dots, a_n) \in K^n : f_i(a_1, \dots, a_n) = 0, 1 \leq i \leq m\}.$$

If $X = \mathcal{V}(f_1, \dots, f_m)$, one also says that f_1, \dots, f_m define X .

A set of the form $\mathcal{V}(f_1, \dots, f_m)$ is called a *closed algebraic set* in K^n . For the moment, we shall only be talking about closed algebraic sets here, and so we usually omit the word "closed."

We write $\text{MaxSpec}(R)$ for the space of maximal ideals of the ring R . Then $\text{MaxSpec}(R) \subseteq \text{Spec}(R)$ and has an inherited Zariski topology. Let K be a field, and let $R = K[x_1, \dots, x_n]$ be a polynomial rings over K . Then there is an injective map $\theta : K^n \rightarrow \text{MaxSpec}(R)$ that sends $P = (\lambda_1, \dots, \lambda_n) \in K^n$ to the maximal ideal of polynomials that vanish at P , which is the kernel of the K -algebra surjection $K[x_1, \dots, x_n] \rightarrow K$ such that $f \mapsto f(P)$. This maximal ideal may also be described as the ideal $(x_1 - \lambda_1, \dots, x_n - \lambda_n)R$. The map θ is injective because, given m , the value of $\lambda_i \in K$ for the point P that maps to m is the unique value of $\lambda \in K$ such that $x_i - \lambda_i \in m$. The closed algebraic sets in K^n are simply the closed sets in the inherited Zariski topology from $\text{MaxSpec}(R)$.

Hilbert's Nullstellensatz provides the information that when K is algebraically closed, the map θ is a bijection, and we may identify K^n with $\text{MaxSpec}(R)$.

Over $\mathbb{R}[x, y]$, $\mathcal{V}(x^2 + y^2 - 1) \subseteq \mathbb{R}^2$ is a circle in the plane, while $\mathcal{V}(xy)$ is the union of the coordinate axes. Note that $\mathcal{V}(x, y)$ is just the origin.

For a while we restrict attention to the case where K is an algebraically closed field such as the complex numbers \mathbb{C} . We want to give algebraic sets a dimension in such a way that K^n has dimension n . Thus, the notion of dimension that we develop will generalize the notion of dimension of a vector space.

We shall do this by associating a ring with X , denoted $K[X]$: it is simply the set of functions defined on X that are obtained by restricting a polynomial function on K^n to X . We define the dimension of X to be the same as the Krull dimension of the ring $K[X]$.

We want to mention a result that tends to show that this notion of dimension is a worthwhile one.

First, consider the problem of describing the intersection of two planes in real three-space \mathbb{R}^3 . The planes might be parallel, i.e., not meet at all. But if they do meet in at least one point, they must meet in a line.

More generally, if one has vector spaces V and W over a field K , both subspaces of some larger vector space, then $\dim(V \cap W) = \dim V + \dim W - \dim(V + W)$. If the ambient vector space has dimension n , this leads to the result that $\dim(V \cap W) \geq \dim V + \dim W - n$. In the case of planes in three-space, we see that that dimension of the intersection must be at least $2 + 2 - 3 = 1$.

Over an algebraically closed field, the same result turns out to be true for algebraic sets! Suppose that V and W are algebraic sets in K^n and that they meet in a point $x \in K^n$. We have to be a little bit careful because, unlike vector spaces, algebraic sets in general may be unions of finitely many smaller algebraic sets, which need not all have the same dimension. Algebraic sets which are not finite unions of strictly smaller algebraic sets are called *irreducible*. Each algebraic set is a finite union of irreducible ones in such a way that none can be omitted: these are called *irreducible components*. We define $\dim_x V$ to be the largest dimension of an irreducible component of V that contains x . The theorem we want to mention is that for any algebraic sets V and W in K^n meeting in a point x , $\dim_x(V \cap W) \geq \dim_x V + \dim_x W - n$. This is a beautiful and useful result: it can be thought of as guaranteeing the existence of a solution (or many solutions) of a family of equations.

We next mention one other sort of problem. Given a specific algebraic set $X = V(f_1, \dots, f_m)$, the set J of all polynomials vanishing on it is closed under addition and multiplication by any polynomial — that is, it is an ideal of $K[x_1, \dots, x_n]$. J always contains the ideal I generated by f_1, \dots, f_m . But J may be strictly larger than I : Hilbert's Nullstellensatz tell us that when K is algebraically closed, the ideal J is the radical of I . We want to point out that given a choice of I , say by a listing of generators, it may be *very* hard to tell whether $J = I$, i.e., to tell whether I is radical.

Here is one example of an open question of this sort. Consider the set of pairs of commuting square matrices of size n . Let $M = M_n(K)$ be the set of $n \times n$ matrices over K . Thus,

$$W = \{(A, B) \in M \times M : AB = BA\}.$$

The matrices are given by their $2n^2$ entries, and we may think of this set as a subset of K^{2n^2} . (To make this official, one would have to describe a way to string the entries of the two matrices out on a line.) Then W is an algebraic set defined by n^2 quadratic equations. If $X = (x_{ij})$ is an $n \times n$ matrix of indeterminates and $Y = (y_{ij})$ is another $n \times n$ matrix

of indeterminates, then we may think of the algebraic set W as defined by the vanishing of the entries of the matrix $XY - YX$. These are the n^2 quadratic equations.

Is the ideal of all functions that vanish on W generated by the entries of $XY - YX$? This is a long standing open question. It is known if $n \leq 3$.

We mention one more very natural but very difficult question about algebraic sets. Suppose that one has an algebraic set $X = V(f_1, \dots, f_m)$. What is the least number of elements needed to define X ? In other words, what is the least positive integer k such that $X = V(g_1, \dots, g_k)$?

Here is a completely specific example. Suppose that we work in the polynomial ring in 6 variables $x_1, \dots, x_3, y_1, \dots, y_3$ over the complex numbers \mathbb{C} and let X be the algebraic set in \mathbb{C}^6 defined by the vanishing of the 2×2 subdeterminants or *minors* of the matrix

$$\begin{pmatrix} x_1 & x_2 & x_3 \\ y_1 & y_2 & y_3 \end{pmatrix},$$

that is, $X = V(f, g, h)$ where $f = x_1y_2 - x_2y_1$, $g = x_1y_3 - x_3y_1$, and $h = x_2y_3 - x_3y_2$. We can think of points of X as representing 2×3 matrices whose rank is at most 1: the vanishing of these equations is precisely the condition for the two rows of the matrix to be linearly dependent. Obviously, X can be defined by 3 equations. Can it be defined by 2 equations? No algorithm is known for settling questions of this sort, and many are open, even for relatively small specific examples. In the example considered here, it turns out that 3 equations are needed. I do not know an elementary proof of this fact — perhaps you can find one!

We next want to give another strong form of Hilbert's Nullstellensatz. Fix an algebraically closed field and fix n , and let $R = K[x_1, \dots, x_n]$ be a polynomial ring. For every set of polynomials $S \subseteq K[x_1, \dots, x_n]$, $\mathcal{V}(S) = \mathcal{V}(I)$, where I is the ideal generated by S , and $\mathcal{V}(I) = \mathcal{V}(\text{Rad } I)$, since $\mathcal{V}(f^n) = \mathcal{V}(f)$, always. Since every ideal is finitely generated, we may choose finitely many elements f_1, \dots, f_m that generate I , or any ideal with the same radical as I , and then $\mathcal{V}(S) = \mathcal{V}(f_1, \dots, f_m) = \mathcal{V}(f_1) \cap \dots \cap \mathcal{V}(f_m)$. We are now ready to prove another strong form of Hilbert's Nullstellensatz. If X is any subset of K^n , we write $\mathcal{I}(X) = \{f \in K[x_1, \dots, x_n] : \text{for all } x \in X, f(x) = 0\}$. Note that if $X = \{x\}$ has one point, then $\mathcal{I}(\{x\}) = m_x$, the maximal ideal consisting of all functions that vanish at x . Also note that $\mathcal{I}(X) = \bigcap_{x \in X} m_x$, and is always a radical ideal. These statements are all valid even without the assumption that K is algebraically closed. When K is algebraically closed, we can also state the following:

Theorem (Hilbert's Nullstellensatz, second strong form). *Let K be an algebraically closed field, and consider the polynomial ring $R = K[x_1, \dots, x_n]$ and algebraic sets in K^n . The functions \mathcal{V} and \mathcal{I} give a bijective order-reversing correspondence between radical ideals of R and closed algebraic sets in K^n .*

Proof. Let I be a radical ideal. We may write $I = (f_1, \dots, f_m)R$ for suitable f_j . We must show that $\mathcal{I}(\mathcal{V}(I)) = I$. The left hand side consists of all polynomials that vanish everywhere that the f_i vanish, and the earlier strong form of Hilbert's Nullstellensatz that

we proved says precisely that if g vanishes on $\mathcal{V}(f_1, \dots, f_m)$, then $g \in \text{Rad}(f_1, \dots, f_m) = (f_1, \dots, f_m)$ in this case, since we assumed that $I = (f_1, \dots, f_m)$ is radical.

What remains to be shown is that if X is an algebraic set then $\mathcal{V}(\mathcal{I}(X)) = X$. But since X is an algebraic set, we have that $X = \mathcal{V}(I)$ for some radical ideal I . Consequently, $\mathcal{V}(\mathcal{I}(X)) = \mathcal{V}(\mathcal{I}(\mathcal{V}(I))) = \mathcal{V}(I)$, since $\mathcal{I}(\mathcal{V}(I)) = I$, by what we proved just above, and $\mathcal{V}(I) = X$. \square

Proposition. *In $X = \text{Spec}(R)$ where R is Noetherian, every closed set Z has finitely many maximal closed irreducible subsets, and it is the union of these. This union is irredundant, i.e., none of the maximal closed irreducible sets can be omitted. The maximal closed irreducible subsets of Z are the same as the maximal irreducible subsets of Z .*

If K is an algebraically closed field, the same statements apply to the closed algebraic sets in K^n .

Proof. The maximal irreducible closed subsets of Z correspond to the minimal primes P_1, \dots, P_n of the radical ideal I such that $V(I) = Z$, and this shows that Z is the union of the maximal irreducible closed sets $Z_i = V(P_i)$ contained in Z .

On the other hand, if Z is a finite union of mutually incomparable irreducible closed sets Z_i , then every irreducible subset W of Z is contained in one of them, for W is the union of the closed subsets $W \cap Z_i$, and so we must have $W = W \cap Z_i$ for some i , and thus $W \subseteq Z_i$. This proves that the Z_i are maximal irreducible subsets, and that none of them can be omitted from the union: if Z_j could be omitted it would be contained in the union of the others and therefore contained in one of the others.

The proof for the case of algebraic sets in K^n is the same. \square

In both contexts, the maximal irreducible closed subsets in Z are called the *irreducible components* of Z .

Irreducible closed algebraic sets in K^n , when K is algebraically closed, are called *algebraic varieties*. (To be precise, they are called *affine algebraic varieties*, but we shall not be dealing in this course with the other kinds. These include the irreducible closed algebraic sets in a projective space over K , which are called *projective varieties*, irreducible open sets in an affine variety, which are called *quasi-affine varieties*, and irreducible open sets in a projective variety, which are called *quasi-projective varieties*. The last type includes the others already mentioned. There is also an abstract notion of variety which is more general, but the most important examples are quasi-projective.)

The notation \mathbb{A}_K^n is used for K^n to emphasize that is being thought of as an algebraic set (rather than as, say, a vector space).

Examples. In \mathbb{A}_K^2 , $\mathcal{V}(x_1x_2) = \mathcal{V}(x_1) \cup \mathcal{V}(x_2)$ gives the representation of the algebraic set which is the union of the axes as an irredundant union of irreducible algebraic sets. This corresponds to the fact that in $K[x, y]$, $(xy) = (x) \cap (y)$. Now consider \mathbb{A}_K^6 where the variables are $x_1, x_2, x_3, y_1, y_2, y_3$, so that our polynomial ring is $R = K[x_1, x_2, x_3, y_1, y_2, y_3]$. Instead of thinking of algebraic sets as lying in \mathbb{A}_K^6 , we shall think instead of them as sets of 2×3 matrices, where the values of the variables x_i and y_j are used to create a matrix as

shown: $\begin{pmatrix} x_1 & x_2 & x_3 \\ y_1 & y_2 & y_3 \end{pmatrix}$. Let $\Delta_1 = x_2y_3 - x_3y_2$, $\Delta_2 = x_1y_3 - x_3y_1$ and $\Delta_3 = x_1y_2 - x_2y_1$ be the three 2×2 minors of this matrix. Consider the algebraic set $\mathcal{V}(\Delta_2, \Delta_3)$. We may think of this as the algebraic set of 2×3 matrices such that the minor formed from the first two columns and the minor formed from the first and third columns vanish. If a matrix is in this set, there are two possibilities. One is that the first column is zero, in which case the two minors involved do vanish. The second case is that the first column is not zero. In this case, the second and third columns are multiples of the first column, and this implies that Δ_1 vanishes. From this we obtain that $\mathcal{V}(\Delta_2, \Delta_3) = \mathcal{V}(x_1, y_1) \cup \mathcal{V}(\Delta_1, \Delta_2, \Delta_3)$. This does turn out to be the decomposition of $\mathcal{V}(\Delta_2, \Delta_3)$ as an irredundant union of irreducible components. The hardest part here is to show that $\mathcal{V}(\Delta_1, \Delta_2, \Delta_3)$ is irreducible.

A topological space is called *Noetherian* if it satisfies DCC on closed sets. Thus, $\text{Spec}(R)$ is Noetherian iff the radical ideals of R have ACC, which is, of course true if R is Noetherian.

Proposition. *A subspace Y of a Noetherian topological space X is Noetherian. A Noetherian space is quasi-compact. A topological space X is Noetherian if and only if every open subspace is quasi-compact, in which case every subspace is quasi-compact. In a Noetherian topological space, every closed subset is the finite irredundant union of its maximal closed irreducible subsets, which are the same as its irreducible subsets.*

Proof. For the first statement, it suffices to show that a non-increasing sequence of closed sets Y_i in Y is stable, and we can write $Y_i = Z_i \cap Y$, where Z_i is closed in X . Then the sequence $Z_1, Z_1 \cap Z_2, \dots, Z_1 \cap \dots \cap Z_n, \dots$ is eventually stable in X , and the intersection of the n th term with Y is $Y_1 \cap \dots \cap Y_n = Y_n$.

Consider next a family of closed sets in X with FIP. We must show the intersection is non-empty. We may assume without loss of generality that the family is closed under intersection. But it has a minimal element, and this must be contained in all of the sets, or we could intersect further, contradicting minimality.

Clearly, if X is Noetherian, then every subset is Noetherian and hence quasi-compact, and so is every open subset. It suffices to show that if every open subset is quasi-compact, then X is Noetherian. If not, let $Z_1 \supset Z_2 \supset \dots \supset Z_n \supset \dots$ be a strictly decreasing sequence of closed sets. Call the intersection Z . Then $X - Z$ is open, and is the strictly increasing union of the open sets $X - Z_n$. This gives an open cover with no finite sub-cover, contradicting the quasi-compactness of X .

Finally, let Z be any closed set in X . If it is not a finite union of irreducibles, take a minimal counter-example. If Z itself is irreducible, we are done. If not then $Z = Z_1 \cup Z_2$, where these are proper closed subsets, and hence each is a finite union of irreducibles, since Z is a minimal counterexample. Once we have Z as a finite union of irreducibles, we can omit terms until we have Z as an irredundant finite union of irreducibles, say $Z = Z_1 \cup \dots \cup Z_n$. Now, if Y is an irreducible set contained in Z , it must be contained in one of Z_i , since it is the union of its intersections with the Z_i , which shows that the Z_i are the maximal irreducible sets contained in Z , as well as the maximal irreducible closed sets contained in Z . \square

The category of closed algebraic sets

We next want to make the closed algebraic sets over an algebraically closed field K into a category. Suppose we are given $X \subseteq K^n$ and $Y \subseteq K^m$. We could write \mathbb{A}_K^n instead of K^n and \mathbb{A}_K^m instead of K^m . We define a function $f : X \rightarrow Y$ to be *regular* if there exist polynomials $g_1, \dots, g_m \in K[x_1, \dots, x_n]$ such that for all points $x \in X$, $f(x) = (g_1(x), \dots, g_m(x))$. Thus, the function f can be given by a polynomial formula in the coordinates. It is easy to verify that the identity function is regular and that the composition of two regular functions is regular. The closed algebraic sets over K become a category if we define $\text{Mor}(X, Y)$ to be the set of regular functions from X to Y .

It may seem a bit artificial to require that a map of $X \subseteq \mathbb{A}_K^n$ to $Y \subseteq \mathbb{A}_K^m$ be induced by a map from \mathbb{A}_K^n to \mathbb{A}_K^m (the polynomials g_j in the definition of regular map actually give a map $K^n \rightarrow K^m$ that happens to take X into Y). However, this is not much different from the situation in topology.

Most of the objects of interest in topology (compact manifolds or compact manifolds with boundary) are embeddable as closed sets in \mathbb{R}^n for some n . If $X \subseteq \mathbb{R}^n$ and $Y \subseteq \mathbb{R}^m$, then every continuous function from X to Y is the restriction of a continuous function from $\mathbb{R}^n \rightarrow \mathbb{R}^m$. To see this, think about the composition $X \rightarrow Y \subseteq \mathbb{R}^m$. The function $X \rightarrow \mathbb{R}^m$ is given by an m -tuple of continuous functions from X to \mathbb{R} . But a continuous function from a closed set $X \subseteq \mathbb{R}^n$ to \mathbb{R} does extend to a continuous function from \mathbb{R}^n to \mathbb{R} : this is the Tietze extension theorem, and uses only that \mathbb{R}^n is a normal topological space.

We now enlarge the category of algebraic sets slightly. Given an algebraic set X and mutually inverse set bijections $\alpha : X' \rightarrow X$ and $\beta : X \rightarrow X'$ we shall think of these maps as giving X' the structure of an algebraic set. We define a map $f : X' \rightarrow Y$ to be *regular* if $f \circ \beta$ is regular, and a map $g : Y \rightarrow X'$ to be *regular* if $\alpha \circ g$ is regular.

Of course if we have also given, say, Y' , the structure of an algebraic set via mutually inverse set isomorphisms $\gamma : Y' \rightarrow Y$ and $\delta : Y \rightarrow Y'$ with an algebraic set Y , then $f : X' \rightarrow Y'$ is *regular* if $\gamma \circ f \circ \beta$ is a regular function from X to Y , while $g : Y' \rightarrow X'$ is *regular* if $\alpha \circ g \circ \delta$ is a regular function from Y to X .

More generally, given any category in which the objects have underlying sets and the morphisms are functions on the underlying sets with, possibly, some further restrictive property (groups and group homomorphisms, rings and ring homomorphisms, and topological spaces and continuous maps are examples), one can make an entirely similar construction: given a bijection $\alpha : X' \rightarrow X$ one can introduce an object with underlying set X' into the category in such a way that α is an isomorphism of that new object with X . In the case of rings, one uses the bijection to introduce addition and multiplication on X' : one adds elements of X' by taking the images of the elements in X , adding them in X , and then applying the inverse bijection to the sum to get an element of X' . One introduces multiplication in X' in an entirely similar way.

Given a closed algebraic set $X \subseteq \mathbb{A}_K^n$, the regular functions to K (i.e., to \mathbb{A}_K^1) have the structure of a K -algebra: the restrictions of polynomials g_1 and g_2 to X have a sum

(respectively, a product) that is regular because it is the restriction of $g_1 + g_2$ (respectively, $g_1 g_2$). This ring is called the *coordinate ring* of X and is denoted $K[X]$. It is a reduced finitely generated K -algebra: if a power of a function is 0, all of its values are nilpotent in K and therefore 0 in K , so that the function is identically zero. The coordinate ring is generated over K by the images of the n functions represented by the variables x_1, \dots, x_n . The function x_i assigns to a point in X its i th coordinate, and so the functions x_i are referred to as *coordinate functions*, which explains why the K -algebra they generate is called the coordinate ring.

$K[X]$ is a homomorphic image of $K[x_1, \dots, x_n]$ under the K -algebra homomorphism that sends the function given by a polynomial $g \in K[x_1, \dots, x_n]$ to its restriction to X . The kernel of this K -algebra homomorphism is the ideal $\mathcal{I}(X)$ of all polynomial functions that vanish on X , and so we have a K -algebra isomorphism $K[x_1, \dots, x_n]/\mathcal{I}(X) \cong K[X]$.

In fact, $\text{Mor}(_, \mathbb{A}_K^1)$ is a contravariant functor from algebraic sets to reduced finitely generated K -algebras. Given a map of algebraic sets $f : X \rightarrow Y$ there is a K -algebra homomorphism $f^* : K[Y] \rightarrow K[X]$ induced by composition; for each $g : Y \rightarrow \mathbb{A}_K^1$, we let $f^*(g) = g \circ f : X \rightarrow \mathbb{A}_K^1$.

Now consider the functor $\text{Hom}_{K\text{-alg}}(_, K)$ from reduced finitely generated K -algebras to algebraic sets. Here the subscript indicates that we are dealing with K -algebra homomorphisms. For this to make sense, we have to give $\text{Hom}_{K\text{-alg}}(R, K)$ the structure of an algebraic set: we do this by choosing a finite set of algebra generators *vecrn* for R over K , and then mapping $\text{Hom}_{K\text{-alg}}(R, K)$ to \mathbb{A}_K^n by sending $\phi \in \text{Hom}_{K\text{-alg}}(R, K)$ to the n -tuple $(\phi(r_1), \dots, \phi(r_n)) \in \mathbb{A}_K^n$. We shall see below that the set of values of this map is a closed algebraic set in \mathbb{A}_K^n , and that, up to isomorphism, this algebraic set is independent of the choice of a finite set of generators for R over K . Thus, $\text{Hom}_{K\text{-alg}}(R, K)$ has the structure of an algebraic set. Moreover, $\text{Hom}_{K\text{-alg}}(_, K)$ is a contravariant functor: if $h : R \rightarrow S$ is a K -algebra homomorphism, we get a map $h^* : \text{Hom}_{K\text{-alg}}(S, K)$ to $\text{Hom}_{K\text{-alg}}(R, K)$ induced by composition: $h^*(\theta) = \theta \circ h$. We shall see that this makes $\mathcal{F} = \text{Hom}_{K\text{-alg}}(_, K)$ into a contravariant functor from reduced finitely generated K -algebras to closed algebraic sets over K .

Note that the elements of $\text{Hom}_{K\text{-alg}}(R, K)$ correspond bijectively with the maximal ideals of R : the maximal ideal is recovered from a given homomorphism as its kernel. On the other hand, we have already seen that for any maximal ideal m , $K \rightarrow R/m$ is an isomorphism μ when K is algebraically closed, and we may compose $R \rightarrow R/m$ with μ^{-1} to obtain a K -algebra homomorphism $R \rightarrow K$ whose kernel is the specified maximal ideal m . Note that if we have $\theta : S \rightarrow K$ and we compose with $f : R \rightarrow S$, the kernel of the composition $R \rightarrow S \rightarrow K$ is the same as the contraction of the kernel of θ to R . Thus, the functor MaxSpec is isomorphic with $\mathcal{G} = \text{Hom}_{K\text{-alg}}(_, K)$, and so we could have worked with this functor instead of \mathcal{G} . In particular, we can give every $\text{MaxSpec}(R)$ the structure of an algebraic set.

Our main result in this direction is:

Theorem. *The procedure for giving $\text{Hom}_{K\text{-alg}}(R, K)$ the structure of an algebraic set described above does produce a bijection with an algebraic set, and changing the choice*

of the finite set of generators for R produces an isomorphic algebraic set. \mathcal{F} and \mathcal{G} as described above are contravariant functors such that $\mathcal{F} \circ \mathcal{G}$ is isomorphic with the identity functor on closed algebraic sets over K , and $\mathcal{G} \circ \mathcal{F}$ is isomorphic with the identity functor on reduced finitely generated K -algebras. Thus, the category of closed algebraic sets and regular functions over the algebraically closed field K is anti-equivalent to the category of reduced finitely generated K -algebras.

Proof. We first note that the points of the closed algebraic set X correspond bijectively in an obvious way with the elements of $\text{Hom}_{K\text{-alg}}(K[X], K)$, and, likewise, with the maximal ideals of $K[X]$. Think of $K[X]$, as usual, as $K[x_1, \dots, x_n]/\mathcal{I}(X)$. The maximal ideals of this ring correspond to maximal ideals of $K[x_1, \dots, x_n]$ containing $\mathcal{I}(X)$. Each such maximal ideal has the form m_y for some $y \in \mathbb{A}_K^n$, and the condition that y must satisfy is that $\mathcal{I}(X) \subseteq m_y$, i.e., that all functions in $\mathcal{I}(X)$ vanish at y , which says that $y \in \mathcal{V}(\mathcal{I}(X))$. By our second strong version of Hilbert's Nullstellensatz, $\mathcal{V}(\mathcal{I}(X)) = X$.

We next note that our procedure for assigning the structure of an algebraic set to $\text{Hom}_{K\text{-alg}}(R, K)$ really does give an algebraic set, which is independent, up to isomorphism, of the choice of the set of generators of R as a K -algebra. To see this, let r_1, \dots, r_n be one set of generators of R . Map $K[x_1, \dots, x_n] \rightarrow R$ using the unique K -algebra homomorphism that sends $x_i \mapsto r_i$, $1 \leq i \leq n$. Let I be the radical ideal which is the kernel of this homomorphism, so that $R \cong K[x_1, \dots, x_n]/I$. The set we assigned to $\text{Hom}_{K\text{-alg}}(R, K)$ is $\{(h(r_1), \dots, h(r_n)) : h \in \text{Hom}_{K\text{-alg}}(R, K)\}$. Each K -homomorphism h is uniquely determined by its values on the generators r_1, \dots, r_n . An n -tuple $(\lambda_1, \dots, \lambda_n)$ can be used to define a K -homomorphism if and only if the elements of I vanish on $(\lambda_1, \dots, \lambda_n)$, i.e., if and only if $(\lambda_1, \dots, \lambda_n) \in \mathcal{V}(I)$. This shows that our map from $\text{Hom}_{K\text{-alg}}(R, K)$ to K^n gives a bijection of $\text{Hom}_{K\text{-alg}}(R, K)$ with the algebraic set $\mathcal{V}(I)$.

Now suppose that r'_1, \dots, r'_m are additional elements of R . For every r'_j we can choose $g_j \in K[x_1, \dots, x_n]$ such that $r'_j = g_j(r_1, \dots, r_n)$. The new algebraic set that we get by evaluating every element $h \in \text{Hom}_{K\text{-alg}}(R, K)$ on $r_1, \dots, r_n, r'_1, \dots, r'_m$ is precisely $X' = \{(\lambda_1, \dots, \lambda_n, g_1(\underline{\lambda}), \dots, g_m(\underline{\lambda})) : \underline{\lambda} \in X\}$, where $\underline{\lambda} = (\lambda_1, \dots, \lambda_n)$. The map $X \rightarrow X'$ that sends $\underline{\lambda} = (\lambda_1, \dots, \lambda_n)$ to $(\lambda_1, \dots, \lambda_n, g_1(\underline{\lambda}), \dots, g_m(\underline{\lambda}))$ is given in coordinates by the polynomials $x_1, \dots, x_n, g_1, \dots, g_m$, and so is a morphism in the category of algebraic sets. Likewise, the map $X' \rightarrow X$ which is simply projection on the first n coordinates is given by polynomials in the coordinates, and these are mutually inverse morphisms of algebraic sets. Thus, $X \cong X'$, as required.

This handles the case where one set of generators is contained in another. But now, if r_1, \dots, r_n and r'_1, \dots, r'_m are two sets of generators, we may compare the algebraic set given by r_1, \dots, r_n with that given by $r_1, \dots, r_n, r'_1, \dots, r'_m$, and then the latter with the algebraic set given by r'_1, \dots, r'_m . This completes the proof of the independence of the algebraic set structure that we are assigning to $\text{Hom}_{K\text{-alg}}(R, K)$ from the choice of K -algebra generators for R .

If $R = K[X]$ and we choose as generators r_i the restrictions of the coordinate functions x_i to R , then the algebraic set we get from $\text{Hom}_{K\text{-alg}}(K[X], K)$ is X itself, and this is the

same identification of X with $\text{Hom}_{K\text{-alg}}(K[X], K)$ that we made in the first paragraph. Thus, if we let $S_X : X \rightarrow \text{Hom}_{K\text{-alg}}(K[X], K)$ as in that paragraph, we get an isomorphism of algebraic sets, for we may use the restricted coordinate functions as the generators to place the algebraic set structure on $\text{Hom}_{K\text{-alg}}(K[X], K) = (\mathcal{G} \circ \mathcal{F})(X)$. We claim that S_X is a natural transformation from the identity functor on the category of algebraic sets over K to $\mathcal{G} \circ \mathcal{F}$. We need to see that if $\theta : X \rightarrow Y$ is a morphism of algebraic sets, then $(\mathcal{G} \circ \mathcal{F})(\theta)$ is the same as θ once we identify $\text{Hom}_{K\text{-alg}}(K[X], K)$ with X and $\text{Hom}_{K\text{-alg}}(K[X], K)$ with Y . Let ϕ_x (resp., ϕ'_y) denote evaluation as at $x \in X$ (resp., $y \in Y$). We need to show that $((\mathcal{G} \circ \mathcal{F})(\theta))(\phi_x) = \phi'_{\theta(x)}$ for all $x \in X$. Now, $\mathcal{F}(\theta)$ acting on $v \in K[Y]$ is $v \circ \theta$, and \mathcal{G} applied to $\mathcal{F}(\theta)$ acts by composition as well, so that its value on ϕ_x is the map that sends $v \in K[Y]$ to $(v \circ \theta)(x) = v(\theta(x))$, which is evaluation at $\theta(x)$, as required.

Finally, we need to see that $\mathcal{F} \circ \mathcal{G}$ is isomorphic to the identity functor on finitely generated reduced K -algebras. The map sends R to $K[\text{Hom}_{K\text{-alg}}(R, K)]$ where $\text{Hom}_{K\text{-alg}}(R, K)$ is viewed as a closed algebraic set as discussed above. Each element r of R maps to a function f_r on the set $\text{Hom}_{K\text{-alg}}(R, K)$ by the rule $f_r(u) = u(r)$. It is immediate that this is a K -algebra homomorphism: call it T_R . We shall show that the T_R give an isomorphism of the identity functor with $\mathcal{F} \circ \mathcal{G}$. We first need to show that every T_R is an isomorphism. We use the fact that $R \cong K[x_1, \dots, x_n]/I$ for some radical ideal I , with the coordinate functions as generators, and it suffices to consider the case where $R = K[x_1, \dots, x_n]/I$. This identifies $\text{Hom}_{K\text{-alg}}(R, K)$ with $\mathcal{V}(I)$, and the needed isomorphism follows from the fact that $K[\mathcal{V}(I)] \cong K[x_1, \dots, x_n]/\mathcal{I}(\mathcal{V}(I)) = K[x_1, \dots, x_n]/I$, again by the second strong version of Hilbert's Nullstellensatz.

The last step is to check that T is a natural transformation. Consider a K -algebra homomorphism $\alpha : R \rightarrow S$. Choose a K -algebra homomorphism γ of polynomial ring $A = K[y_1, \dots, y_m]$ onto R with kernel I and a K -algebra homomorphism δ of a polynomial ring $B = K[x_1, \dots, x_n]$ onto S with kernel J . Without loss of generality, we may assume that $R = A/I$, $S = B/J$. Choose $g_1, \dots, g_m \in K[x_1, \dots, x_n]$ such that the image of y_j in R maps to the image of g_j in B , $1 \leq j \leq m$, so that α is induced by the K -algebra map $A \rightarrow B$ that sends y_j to g_j , $1 \leq j \leq m$. The corresponding map of algebraic sets $\mathcal{V}(J) \rightarrow \mathcal{V}(I)$ is given in coordinates by the g_j . Finally, the induced map $K[\mathcal{V}(I)] \cong A/\mathcal{I}(\mathcal{V}(I)) = A/I$ to $K[\mathcal{V}(J)] \cong B/\mathcal{I}(\mathcal{V}(J)) = B/J$ is induced by composition with the map given by the polynomials g_1, \dots, g_m . This means that the image of an element of A/I represented by $P(y_1, \dots, y_m) \in A$ is represented by the coset in B/J of $P(g_1, \dots, g_m) \in B$, and this shows that with the identifications we are making, $\mathcal{F} \circ \mathcal{G}(\alpha)$ is α , which is exactly what we need. \square

Given an algebraic set X over an algebraically closed field K , we define $\dim(X)$ to be the same as $\dim(K[X])$. The dimension of a ring is the supremum of the dimensions of its quotients by minimal primes. Thus, $\dim(X)$ is the same as the supremum of the dimensions of the irreducible components of X . Evidently, $\dim(X)$ is also the same as the supremum of lengths of chains of irreducible closed subsets of X . We define the dimension of X near a point $x \in X$ to be the supremum of the dimensions of the irreducible components of X that contain x . If the corresponding maximal ideal of $R = K[X]$ is $m = m_x$, this is also the dimension of R_m : it has minimal primes P corresponding precisely to the irreducible

components $V(P)$ that contain x , and the length of any saturated chain from P to $m = \dim(R_m/PR_m) = \dim(R/P) =$ the dimension of the irreducible component $V(P)$, from which the result follows.

There are at least three ways to think of an algebra R over a commutative ring K . It is worth considering all three points of view. One is purely algebraic: R is an abstract algebraic environment in which one may perform certain sorts of algebraic manipulations.

A second point of view is to think of R , or rather some topological space associated with R , as a geometric object. We have seen explicitly how to do this when R is a finitely generated reduced K -algebra and K is an algebraically closed field. But a geometric point of view, introduced by A. Grothendieck, can be taken in great generality, when R is *any* commutative ring. In Grothendieck's theory of schemes, a geometric object $\text{Spec}(R)$, is introduced that has more structure than just the topological space of prime ideals of R that we have talked about here. The geometric point of view has been very effective as a tool in commutative algebra, even if one is only interested in seemingly purely algebraic properties of rings.

The third point of view is simplest when R is a finitely generated algebra over a Noetherian ring K (and it is simplest of all when K is a field). In this case one has that $R = K[x_1, \dots, x_n]/(f_1, \dots, f_m)$. Now let S be any K -algebra. Then $\text{Hom}_{K\text{-alg}}(R, S)$ is in bijective correspondence with the set of solutions of the set of m simultaneous equations

$$\begin{array}{c}
 f_1(x_1, \dots, x_n) = 0 \\
 \dots \\
 (*) \quad \dots \\
 \dots \\
 f_m(x_1, \dots, x_n) = 0
 \end{array}$$

in S^n , for to give a K -homomorphism from R to S is the same as to give an n -tuple of elements of S (which will serve as the values of the homomorphism on the images of the variables x_1, \dots, x_n) that satisfy these equations. The set of homomorphisms $\text{Hom}_{K\text{-alg}}(R, S)$ is called the *set of S -valued points* of the scheme $\text{Spec}(R)$ in scheme theory: since we don't have that theory available, we shall simply refer to it as the set of S -valued points of R . Recall again that K can be any Noetherian ring here. This point of view can be extended: we do not need to assume that R is finitely generated over K , nor that K is Noetherian, if we allow infinitely many variables in our polynomial ring, and infinite families of polynomial equations to solve. Thus, very generally, a K -algebra may be thought of as an encoded system of equations. When one takes homomorphisms into S , one is solving the equations in S . A different way to say this is the following: suppose that we start with a system of equations over K , and define a functor from K -algebras to sets that assigns to every K -algebra S the set of solutions of the family of equations such that the values of the variables are in S . If one forms the polynomial ring in the variables occurring and then the quotient by the ideal generated by the polynomials set equal to 0 in the equations, the resulting K -algebra represents this functor.

Here is an example. Let $B = \mathbb{R}[X, Y, Z]/(X^2 + Y^2 + Z^2 - 1) = \mathbb{R}[x, y, z]$, and let $S = B[U, V, W]/(xU + yV + zW) = \mathbb{R}[x, y, z, u, v, w]$. We can also form B in a single

step as $\mathbb{R}[X, Y, Z, U, V, W]/(X^2 + Y^2 + Z^2 - 1, XU + YV + ZW)$. The \mathbb{R} -homomorphisms from B or \mathbb{R} -valued points of B correspond to the set $\{(a, b, c) \in \mathbb{R}^3 : a^2 + b^2 + c^2 = 1\}$: the real 2-sphere of radius one centered at the origin in \mathbb{R}^3 . The \mathbb{R} -valued points of S correspond to pairs $(a, b, c), (d, e, f)$ such that $(a, b, c) \in S^2$ and $(a, b, c) \cdot (d, e, f) = 0$, which means that the vector (d, e, f) represents a tangent vector to the sphere at the point (a, b, c) . That is, the \mathbb{R} -valued points of S correspond to the points of the tangent bundle to the real 2-sphere.

Products

Let K be an algebraically closed field. Given two algebraic sets $X = \mathcal{V}(I) \subseteq K^m = \mathbb{A}_K^m$, where we use x_1, \dots, x_m for coordinates, and $Y = \mathcal{V}(J) \subseteq K^n = \mathbb{A}_K^n$, where we use y_1, \dots, y_n for coordinates, the set $X \times Y \subseteq K^{m+n} = \mathbb{A}_K^{m+n}$ is an algebraic set defined by the expansions of I and J to $K[x_1, \dots, x_m, y_1, \dots, y_n] \cong K[x_1, \dots, x_m] \otimes_K K[y_1, \dots, y_n]$. It is obvious that a point satisfies both the conditions imposed by the vanishing of I and of J if and only if its first m coordinates give a point of X and its last n coordinates give a point of Y .

Let $S = K[x_1, \dots, x_m]$ thought of as $K[\mathbb{A}_K^m]$ and $T = K[y_1, \dots, y_n]$ thought of as $K[\mathbb{A}_K^n]$. Then

$$K[X \times Y] \cong (S \otimes_K T) / \text{Rad}(I^e + J^e),$$

where the superscript e indicates expansion of ideals. Since

$$(S \otimes_K T) / (I^e + J^e) \cong (S \otimes_K T) / (I \otimes_K T + S \otimes_K J) \cong (S/I) \otimes_K (T/J),$$

we have that

$$K[X \times Y] \cong ((S/I) \otimes_K (T/J))_{\text{red}} \cong (K[X] \otimes_K K[Y])_{\text{red}}.$$

It is not necessary to kill the nilpotents, because of the following fact:

Theorem. *Let R and S be algebras over an algebraically closed field K .*

- (a) *If R and S are domains, then $R \otimes_K S$ is a domain.*
- (b) *If R and S are reduced, then $R \otimes_K S$ is reduced.*

Proof. For part (a), let \mathcal{F} denote the fraction field of R . Since K is a field, every K -module is free, and, therefore, flat. We have an injection $R \hookrightarrow \mathcal{F}$. Thus, $R \otimes_K S \hookrightarrow \mathcal{F} \otimes_K S$. By Supplementary Problem Set #4, problem 6., this ring is a domain, and so its subring $R \otimes_K S$ is a domain.

For part (b), note that R is the directed union of its finitely generated K -subalgebras R_0 . Thus, $R \otimes_K S$ is the directed union of its subalgebras $R_0 \otimes_K S$ where $R_0 \subseteq R$ is finitely generated. Similarly, this ring is the directed union of its subalgebras $R_0 \otimes_K S_0$, where both $R_0 \subseteq R$ and $S_0 \subseteq S$ are finitely generated. We can therefore reduce to the case where R and S are finitely generated.

Let P_1, \dots, P_m be the minimal primes of R . Since R is reduced, their intersection is 0. Therefore, R injects into $\prod_i (R/P_i)$. Thus,

$$R \otimes_K S \hookrightarrow \left(\prod_i (R/P_i) \right) \otimes_K S \cong \prod_i \left((R/P_i) \otimes_K S \right)$$

(if we think of the products as direct sums, we have an obvious isomorphism of K -vector spaces: the check that multiplication is preserved is straightforward), and so it suffices to show that each factor ring of this product is reduced. Thus, we need only show that if R is a domain and S is reduced, where these are finitely generated K -algebras, then $R \otimes_K S$ is reduced. But now we may repeat this argument using the minimal primes Q_1, \dots, Q_n of S , and so we need only show that each ring $R \otimes_K (S/Q_j)$ is reduced, where now both R and S/Q_j are domains. By part (a), these tensor products are domains. \square

One may also show that the tensor product of two reduced rings over an algebraically closed field is reduced using an equational argument and Hilbert's Nullstellensatz, similar to the argument for Supplementary Problem Set #4, **6**.

We return to the study of algebraic sets over an algebraically closed field. We have now established an isomorphism $K[X \times Y] \cong K[X] \otimes_K K[Y]$. Moreover, it is easy to see that the product projections $X \times Y \rightarrow X$, $X \times Y \rightarrow Y$ correspond to the respective injections $K[X] \rightarrow K[X] \otimes_K K[Y]$ and $K[Y] \rightarrow K[X] \otimes_K K[Y]$, where the first sends $f \mapsto f \otimes 1$ and the second sends $g \mapsto 1 \otimes g$.

From the fact that $K[X] \otimes_K K[Y]$ is a coproduct of $K[X]$ and $K[Y]$ in the category of K -algebras, it follows easily that $X \times Y$ (with the usual product projections) is a product of X and Y in the category of algebraic sets. That is, giving a morphism from Z to $X \times Y$ is equivalent to giving a pair of morphisms, one from Z to X and the other from $Z \rightarrow Y$. This is simply because giving a morphism from Z to $X \times Y$ is equivalent to giving a K -homomorphism $K[X] \otimes_K K[Y]$ to $K[Z]$, which we know is equivalent to giving a K -homomorphism $K[X] \rightarrow K[Z]$ and a K -homomorphism $K[Y] \rightarrow K[Z]$: as already mentioned, $K[X] \otimes_K K[Y]$ is a coproduct for $K[X]$ and $K[Y]$ in the category of K -algebras. Notice also that since $K[X] \otimes_K K[Y]$ is a domain whenever $K[X]$ and $K[Y]$ are both domains, we have:

Corollary. *The product of two varieties (i.e., irreducible algebraic sets) in \mathbb{A}_K^n over an algebraically closed field K is a variety (i.e., irreducible).*

We also note:

Proposition. *If X and Y are algebraic sets over the algebraically closed field K , then*

$$\dim(X \times Y) = \dim(X) + \dim(Y).$$

Proof. $K[X]$ is module-finite over a polynomial ring A in d variables where $d = \dim(X)$, say with module generators u_1, \dots, u_s , and $K[Y]$ is module-finite, say with module generators v_1, \dots, v_t , over a polynomial ring B in d' variables. Hence, $K[X] \otimes_K K[Y]$ is module-finite (with module generators $u_i \otimes v_j$) over a polynomial ring in $d + d'$ variables.

Note that $A \otimes_K B$ injects into $A \otimes_K K[Y]$ because A is K -flat, and the latter injects into $K[X] \otimes_K K[Y]$ because $K[Y]$ is K -flat. \square

The dimension of the intersection of two varieties

We next prove a result that was promised long ago:

Theorem. *Let X and Y be irreducible algebraic sets meeting at a point $x \in \mathbb{A}_K^n$, where K is an algebraic closed field. Then*

$$\dim(X \cap Y) \geq \dim(X) + \dim(Y) - n.$$

In fact every irreducible component of $X \cap Y$ has dimension $\geq \dim(X) + \dim(Y) - n$.

Proof. Let $X = \mathcal{V}(P)$ and $Y = \mathcal{V}(Q)$, where P and Q are prime ideals of $K[x_1, \dots, x_n]$. Then $X \cap Y = \mathcal{V}(P + Q)$, although $P + Q$ need not be radical, and

$$K[X \cap Y] = (K[x_1, \dots, x_n]/(P + Q))_{\text{red}}.$$

Now

$$K[x_1, \dots, x_n]/(P + Q) \cong ((K[x_1, \dots, x_n]/P) \otimes_K (K[y_1, \dots, y_n]/Q'))/I_\Delta,$$

where I_Δ is the ideal generated by the $x_i - y_i$ for $1 \leq i \leq n$, which is the ideal that defines the diagonal Δ in $\mathbb{A}_K^n \times_K \mathbb{A}_K^n$. The point is that once we kill the generators $x_i - y_i$ of I_Δ , the ring $K[y_1, \dots, y_n]$ is identified with $K[x_1, \dots, x_n]$, and the image of Q' is Q . (Geometrically, we are identifying $X \cap Y$ with $(X \times Y) \cap \Delta$ in $\mathbb{A}_K^n \times \mathbb{A}_K^n$, via the map $z \mapsto (z, z)$.) Let $R = (K[x_1, \dots, x_n]/P) \otimes_K (K[y_1, \dots, y_n]/Q')$. The dimension of $R = K[X \times Y]$ is $\dim(X) + \dim(Y)$. Since the intersection $X \cap Y$ is non-empty, we know that I_Δ expands to a proper ideal. The dimension of the quotient will be the supremum of the heights of the m/I_Δ as m runs through maximal ideals containing I_Δ , and this will be the supremum of the dimensions of the local rings $\dim(R_m/I_\Delta R_m)$. Each R_m has dimension equal to that of R , i.e., $\dim(X) + \dim(Y)$. But I_Δ is generated by n elements, and killing n elements in the maximal ideal of a local ring drops the dimension of the local ring by at most n . Thus, every $R_m/I_\Delta R_m$ has dimension at least $\dim(X) + \dim(Y) - n$, and the result follows. To get the final statement, let x be a point of the irreducible component considered not in any other irreducible component of $X \cap Y$, and let m be the corresponding maximal ideal of R . We have that $R_m/I_\Delta R_m$ has dimension at least $\dim(X) + \dim(Y) - n$ as before, but now there is a unique minimal prime P in this ring, corresponding to the fact that only one irreducible component of $X \cap Y$ contains x . It follows that this irreducible component has dimension at least $\dim(X) + \dim(Y) - n$. \square

Note that the argument in the proof shows that the map $X \cap Y \rightarrow (X \times Y) \cap \Delta$ that sends z to (z, z) is an isomorphism of algebraic sets.

Recall that $\dim_x(X)$ is the largest dimension of an irreducible component of X that contains x . It follows at once that:

Corollary. *Let X and Y be algebraic sets in K^n , where K is an algebraically closed field, and suppose $x \in X \cap Y$. Then*

$$\dim_x(X \cap Y) \geq \dim_x(X) + \dim_x(Y) - n.$$

Proof. Let X_0 be an irreducible component of X containing x of largest dimension that contains x and Y_0 be such a component of Y with $x \in Y_0$. Then $\dim_x(X) = \dim(X_0)$ and $\dim_x(Y) = \dim(Y_0)$. Apply the result for the irreducible case to X_0 and Y_0 . \square

The theorem we have just proved may be thought of as an existence theorem for solutions of equations: given two sets of equations in n variables over an algebraically closed field, if the two sets of equations have a common solution x , and the solutions of the first set have dimension d near x while the solutions of the second set have dimension d' near x , then the set of simultaneous solutions of the two sets has dimension at least $d + d' - n$ near x . This is well known for solutions of linear equations, but surprising for algebraic sets!

Open and locally closed algebraic sets

A subset of a topological space is called *locally closed* if it is, equivalently, (1) the intersection of an open set with a closed set, (2) a closed subset of an open set, or (3) an open subset of a closed set. Let $X \subseteq \mathbb{A}_K^n$ be a closed algebraic set. Let $f \in K[X] = R$, and let $X_f = \{x \in X : f(x) \neq 0\}$. Then X_f corresponds bijectively to the set of maximal ideals in R_f . Therefore, X_f has the structure of a closed algebraic set (*a priori*, it is only a locally closed algebraic set). If we think of R as $K[x_1, \dots, x_n]/I$ where $I = \mathcal{I}(X)$, we can map $K[x_1, \dots, x_{n+1}] \rightarrow R_f$, extending the map $K[x_1, \dots, x_n] \rightarrow R$ by mapping $x_{n+1} \rightarrow 1/f$. X now corresponds bijectively to a closed algebraic set in \mathbb{A}_K^{n+1} : the bijection sends x to $(x, 1/f(x))$. The closed algebraic set in question may be described as $\{(x, \lambda) \in \mathbb{A}_K^{n+1} : x \in X \text{ and } \lambda = 1/f(x)\}$. The new defining ideal is $I + (fx_{n+1} - 1)$.

We define a function $X_f \rightarrow K$ to be *regular* if it is regular with respect to the closed algebraic set structure that we have placed on X_f . This raises the following question: suppose that we have a cover of a closed algebraic set X by open sets X_{f_i} and a function $g : X \rightarrow K$ such that the restriction of g to each X_{f_i} is regular in the sense just specified. Is g regular? We shall show that the answer is “yes,” and this shows that regularity is a local property with respect to the Zariski topology. Let g_i denote the restriction of g to $X_i = X_{f_i}$. Note that $g_i|_{X_j} = g_j|_{X_i}$ for all i, j , since they are both restrictions of g .

The following fact gives a generalization to arbitrary modules over an arbitrary commutative ring, and underlies the theory of schemes.

Theorem. *Let R be any ring and M any R -module. Let $X = \text{Spec}(R)$, and let f_i be a family of elements of R such that the open sets $X_i = X_{f_i} = D(f_i)$ cover X . Suppose that for every i we are given an element $u_i \in M_{f_i} = M_i$, and suppose that (*) for all choices of i and j , the images of u_i and u_j in $M_{f_i f_j}$ agree. Then there is a unique element $u \in M$ such that for all i , the image of u in M_{f_i} is u_i .*

The result says, informally, that “constructing” an element of a module is a local problem: one can solve it on an open cover, provided the solutions “fit together” on overlaps.

This turns many problems into local problems: for example, if M is finitely presented, the problem of constructing a map of modules $M \rightarrow N$ amounts to giving an element of the module $\text{Hom}_R(M, N)$. Since localization commutes with Hom when M is finitely presented, the problem of doing the construction becomes local.

Note that if we apply this result in the case of the algebraic set X , we find that there is an element $g_0 \in K[X]$ whose image in $K[X_i]$ is g_i for all i . This implies that g_0 agrees with g on X_i . Since the X_i cover X , $g_0 = g$. Thus, $g \in K[X]$. Consequently, the theorem stated above does show that regularity is a local property.

Proof of the theorem. Uniqueness is obvious: if u and u' are two such elements, then they agree after localizing at any f_i . When one localizes at a prime P , since P cannot contain all the f_i , u and u' have the same image in M_P . It follows that $u = u'$. We focus on the existence of u .

The statement that the X_i cover is equivalent to the statement that the f_i generate the unit ideal. Then finitely many generate the unit ideal: call these f_{i_1}, \dots, f_{i_n} . Suppose that we can construct $u \in M$ such that the image of u is $u_{i_t} \in M_{i_t}$, $1 \leq t \leq n$. We claim that the image u'_j of u in M_j is u_j for any j . To see this, it suffices to show that $u'_j - u_j$ vanishes in $(M_j)_P$ for any $P \in X_j$. But X_j is covered by the sets $X_j \cap X_{i_t}$, $1 \leq t \leq n$. If $P \in X_{i_t}$, it suffices to show that u'_j and u_j have the same image in $M_{f_{i_t}f_j}$. The image of u'_j is the same as the image of u , and hence the same as the image of u_{i_t} , and the result follows from our assumption (*).

Therefore, it suffices to work with the cover by the $X_{f_{i_t}}$, and we simplify notation: we let the index set be $\{1, \dots, n\}$ and so the f_i s are simply f_1, \dots, f_n , the cover is X_1, \dots, X_n , and $M_i = M_{f_i}$. We use induction on n . If $n = 1$, $X_1 = X$ and the result is clear: $u = u_1$.

We next consider the case where $n = 2$. This is the core of the proof. Let $u_1 = v_1/f_1^s$ and $u_2 = v_2/f_2^t$ where $v_1, v_2 \in M$. Since these agree in $M_{f_1f_2}$ there exists an integer N such $f_1^N f_2^N (f_2^t v_1 - f_1^s v_2) = 0$. Then $u_1 = f_1^N v_1 / f_1^{N+s}$, $u_2 = f_2^N v_2 / f_2^{N+t}$, and

$$f_2^{N+t} f_1^N v_1 - f_1^{N+s} f_2^N v_2 = (f_1 f_2)^N (f_2^t v_1 - f_1^s v_2) = 0.$$

Thus, if we replace f_1 by f_1^{N+s} , f_2 by f_2^{N+t} , v_1 by $f_1^N v_1$ and v_2 by $f_2^N v_2$, then $u_1 = v_1/f_1$, $u_2 = v_2/f_2$, and $f_2 v_1 - f_1 v_2 = 0$ (Note that the original f_1^{N+s} and f_2^{N+t} generate the unit ideal, since any maximal ideal containing both would have to contain both f_1 and f_2 , a contradiction: thus, the new f_1 and f_2 still generate the unit ideal).

Choose r_1, r_2 such that $r_1 f_1 + r_2 f_2 = 1$. Let $u = r_1 v_1 + r_2 v_2$. Then

$$f_1 u = r_1 f_1 v_1 + r_2 (f_1 v_2) = r_1 f_1 v_1 + r_2 (f_2 v_1) = (r_1 f_1 + r_2 f_2) v_1 = v_1,$$

so that $u = v_1/f_1$ in M_1 , and $u = v_2/f_2$ in M_2 by symmetry.

We now assume that $n > 2$ and that the result has been established for integers $< n$. Suppose that

$$r_1 f_1 + \dots + r_n f_n = 1.$$

Let

$$g_1 = r_1 f_1 + \cdots + r_{n-1} f_{n-1}$$

and $g_2 = f_n$. Evidently, g_1 and g_2 generate the unit ideal, since $g_1 + r_n g_2 = 1$. Consider the images of f_1, \dots, f_{n-1} in R_{g_1} . Because g_1 is invertible, they generate the unit ideal. We now apply the induction hypothesis to M_{g_1} , using the images of the f_i for $1 \leq i \leq n-1$ to give the open cover of $\text{Spec}(R_{g_1})$. Let u'_i denote the image of u_i in $M_{g_1 f_i}$, $1 \leq i \leq n-1$. It is straightforward to verify that condition (*) continues to hold here, using cases of the original condition (*). By the induction hypothesis, there is an element of M_{g_1} , call it w_1 , such that the image of w_1 in each $M_{g_1 f_i}$ is the same as the image of u_i , $1 \leq i \leq n-1$. We claim that the images of w_1 and u_n agree in $M_{g_1 f_n}$. It suffices to show that they agree after localizing at any prime P , and P cannot contain the images of all of f_1, \dots, f_{n-1} . If P does not contain f_i , $1 \leq i \leq n-1$, the result follows because the images of u_i and u_n agree in $M_{f_i f_n}$. We can now apply the case where $n = 2$ to construct the required element of M . \square

The local nature of the regularity of a morphism

Corollary. *Let X and Y be closed algebraic sets over an algebraically closed field K . Then a function $h : X \rightarrow Y$ is regular if and only if (#) it is continuous and for all $x \in X$ there is an open neighborhood Y_g of $y = h(x)$ and an open neighborhood $X_f \subseteq h^{-1}(y)$ such that the restriction of h mapping X_f to Y_g is regular.*

Proof. $Y \subseteq \mathbb{A}_K^n$ (with coordinates x_1, \dots, x_n in the latter), and we will reduce to showing that the composite map $X \rightarrow \mathbb{A}_K^n$ is regular. Let h_i be the composition of this map with the i th coordinate projection. It suffices to show that every h_i is regular. Let X_f be a neighborhood of $x \in X$ such that h maps into an open neighborhood Y_g of $h(x)$. It will correspond to a K -algebra homomorphism $K[Y]_g \rightarrow K[X]_f$. Note that g is the restriction of a function g' on \mathbb{A}_K^n , and $(\mathbb{A}_K^n)_{g'}$ meets Y in Y_g . The inclusion $Y \subseteq \mathbb{A}_K^n$ corresponds to a surjection $K[x_1, \dots, x_n] \rightarrow K[Y]$. The map $Y_g \rightarrow (\mathbb{A}_K^n)_{g'}$ corresponds to the ring map $K[x_1, \dots, x_n]_{g'} \rightarrow K[Y]_g$ induced by localization at the multiplicative system generated by g' (recall that g' maps to g). Thus, the map $X_f \rightarrow (\mathbb{A}_K^n)_{g'}$ is regular, and so is the map $X_f \rightarrow \mathbb{A}_K^n$, which corresponds to the composite ring map

$$K[x_1, \dots, x_n] \rightarrow K[x_1, \dots, x_n]_{g'} \rightarrow K[Y]_g \rightarrow K[X]_f.$$

It follows that the composition of the map $X_f \rightarrow \mathbb{A}_K^n$ with the i th coordinate projection is regular: this is the restriction of h_i to X_f . Since the X_f cover X , it follows that every h_i is regular, and so h is regular. \square

We can now define when a function between open subsets of algebraic sets (i.e., locally closed algebraic sets) is a morphism: simply use the condition (#) in the Corollary.

How can one generalize further?

The category already described can be embedded, up to isomorphism, in a much larger category. One way to do this is to introduce the category of schemes: this category can be

thought of as containing both the category of algebraic sets over K (reduced quasi-affine schemes over an algebraically closed field K) and the opposite of the category of commutative rings with identity (the category of affine schemes). There are other generalizations as well. Here, we shall only give a glimpse of the much larger categories that one may consider.

A set has the structure of a *reduced scheme of finite type* over an algebraically closed field K if it is a topological space X with a finite open cover by sets X_i together with, for every i , a bijection $f_i : X_i \cong Y_i$ where Y_i is a closed algebraic set over K , satisfying the additional condition that if

$$f_{ij} : X_i \cap X_j \cong f_i(X_i \cap X_j) = Y_{ij} \subseteq Y_i,$$

then the for all i, j the composite

$$f_{ji} \circ f_{ij}^{-1} : Y_{ij} \rightarrow Y_{ji}$$

is an isomorphism of (locally closed) algebraic sets.

Roughly speaking, a reduced scheme of finite type over K is the result of pasting together finitely many closed algebraic sets along open overlaps that are isomorphic in the category of locally closed algebraic sets. This is analogous to the definitions of topological, differentiable and analytic manifolds by pasting open subsets having the same structure as an open set in \mathbb{R}^n (or \mathbb{C}^n in the case of an analytic manifold).

One can use condition (#) to define when a function between two reduced schemes of finite type over K is a morphism: thus, we require that f be continuous, and that for all $x \in X$, if $y = f(x)$, then when we choose an open neighborhood V of y with the structure of a closed algebraic set, and an open neighborhood U of x with the structure of a closed algebraic set such that $f(U) \subseteq V$, then restriction of f to a map from U to V is a morphism of algebraic sets. Our results on the local character of morphisms show that when X and Y are closed algebraic sets, we have not enlarged the set of morphisms from X to Y .

A major failing of this theory is that while the category of finitely generated K -algebras has rings with nilpotents, our reduced schemes never have any. It turns out that the presence of nilpotents can carry geometric information! Even if one detests nilpotents and never wants them around, it is very useful on occasion to be able to say that there really aren't any because of a suitable theorem (as opposed to saying that there aren't any because we were forced by our definitions to kill them all). For example, one cannot express the fact that the tensor product of two reduced K -algebras is reduced in the category of reduced schemes. While there is an object corresponding to the reduced tensor product, there is no object corresponding to the tensor product. The remedy is the theory of schemes: as indicated earlier, the category of schemes contains the opposite of the category of rings as a subcategory, and contains the category of reduced schemes of finite type over an algebraically closed field as well.

When one does the full theory of schemes, the definition of a reduced scheme of finite type over an algebraically closed field K is somewhat different, but the category of reduced

schemes of finite type over K introduced here is equivalent to the category one gets from the more general theory of schemes.