8300 2003 day 2: Hilbert's Nullstellensatz

Motivation: Root fields for polynomials in several variables

We know if k is any field that a polynomial f(X) in k[X] always has a root in some finite algebraic extension field E of k. In fact if g(X) is any irreducible factor of f in the ufd (unique factorization domain) k[X], then E = k[X]/(g) is such an extension field, where x = X mod(g) is a root, since g(X) is zero, mod g(X). E is finite dimensional over k since it is spanned as vector space by the monomials $1,X,X^2,....,X^{d-1}$ of degree less than d = deg(g). Since g divides f, the irreducible polynomial g generates a maximal ideal (g) containing the ideal (f), in the pid (principal ideal domain) k[X].

What about polynomials of several variables? I.e. if $f(X_1,...,X_n)$ is a polynomial in $k[X_1,...,X_n]$, is there a finite algebraic extension field E of k in which f has a "root", (i.e. a coordinate vector $(a_1,...,a_n)$ in E^n at which f vanishes)? More generally, if $\{f_j\}$ is a collection of polynomials in $k[X_1,...,X_n]$, is there a finite algebraic extension of k in which all the f_j have a common zero? This might seem obvious, but it is not so trivial to prove. In fact it is not always true for a simple reason. If a is a common root of all polynomials $\{f_j\}$ then it is also a root of any linear combination of them, i.e. of any polynomial of form $\sum g_jf_j$ with g_j in $k[X_1,...,X_n]$. So if the collection $\{f_j\}$ generates the "unit ideal" in the polynomial ring, i.e. if 1 is a linear combination of the $\{f_j\}$, then there cannot be any common roots. So the correct question to ask is whether every collection $\{f_j\}$ that generates a "proper" ideal (i.e. smaller than the unit ideal) has a common root, or equivalently whether every proper ideal I in $k[X_1,...,X_n]$ has a common root, in some finite algebraic extension field E of k. This is true as we will see, and is called the nullstellensatz, (originally due to Hilbert).

First let's see what such an extension field F would look like. (I will write k[X] sometimes for the polynomial ring in n variables.) For any extension F of k, there is a well defined k - algebra map, "evaluation at a", from k[X]-->F, taking $f(X_1,...,X_n)$ to $f(a_1,...,a_n)$. Moreover, if M is the kernel of this map, then k[X]/M embeds into F, so that we have inclusions k in k[X]/M in F. Since every polynomial in I vanishes at a, we I is contained in M, and since k[X]/M is embedded in the field F, M is a prime ideal. We claim M is a maximal ideal.

Since F is assumed to be a finite algebraic extension of k, hence a finite k vector space, then k[X]/M = E is a domain which is also finite as a vector space over k. Since multiplication by a non zero element u of the domain E gives an injective k linear map of this domain to itself, it must be surjective, hence there is an element v also in E with uv = 1. Thus E is an algebraic field extension of k of form E = k[X]/M, where I is in M = a maximal ideal, and since the common root a of the ideal I equals the value of X (= (X1,...,Xn)) at a, the root a actually lives in E^n.

This shows that the only place to look for common roots of an ideal I in $k[X_1,...,X_n]$, is in quotient fields E of form $k[X_1,...,X_n]/M$, where M is a maximal ideal of $k[X_1,...,X_n]$ containing I, just as in the case of one variable. Thus the question arises, for which maximal ideals M in $k[X_1,...,X_n]$ is the field $k[X_1,...,X_n]/M$ actually finite algebraic over k? We will prove eventually that this is true for all maximal ideals. If you look back at the proof in one variable you will see that it used the division algorithm, something which is notoriously lacking in the elementary theory of polynomials of several variables. I.e. given a finite set of polynomials $\{f_j\}$ in several variables, and another polynomial g, how do you decide in a finite number of steps, by some kind of repeated division, whether g is a linear combination of the polynomials $\{f_j\}$? There is a solution to this question today, in terms of "Grobner bases" for an ideal, and with this theory one can prove the nullstellensatz in a more algorithmic way. We will give a more classical, less constructive proof.

Before proving the general theorem, let's look at a special case where the ideal I has a common zero vector a in the base field k, the field of coefficients of the polynomials in I. Since a = $(a_1,...,a_n)$ where all components a_i belong to k, the evaluation map $k[X_1,...,X_n]$¨k has values in k, and has kernel M a maximal ideal of $k[X_1,...,X_n]$ containing I. Moreover since evaluation takes each variable X_i to a_i, the kernel M contains the polynomials X_i-a_i for all i = 1,...,n. Now the ideal $(X_1-a_1,...,X_n-a_n)$ is already maximal, so must equal M. I.e. in case the ideal I in k[X] has a common zero vector a with components in k, the kernel of the corresponding evaluation map at a has form $(X_1-a_1,...,X_n-a_n)$. Conversely if I is contained in any maximal ideal of form $(X_1-a_1,...,X_n-a_n)$, with all a_i in k, then the map evaluation at a = $(a_1,...,a_n)$ is a surjective k algebra map $k[X_1,...,X_n]$-->k whose kernel contains I, so a is a common zero of I with coefficients in k. Equivalently, if f:$k[X_1,...,X_n]$-->k is any k algebra map with kernel M, and if $f(X_i) = a_i$, then M contains the maximal ideal $(X_1-a_1,...,X_n-a_n)$ hence equals it, and a is a common zero of any ideal contained in ker(f).

More simply, the arguments above show there are one one correspondences between the following sets for any field k:

{points $(a_1,...,a_n)$ in k^n}

≈ {k algebra maps f:$k[X_1,...,X_n]$-->k}

≈ {maximal ideals M in $k[X_1,...,X_n]$ such that k-->$k[X_1,...,X_n]/M$ is an isomorphism}

≈ {maximal ideals of form $(X_1-a_1,...,X_n-a_n)$ in $k[X_1,...,X_n]$}

These correspondences are as follows: a point a in k^n yields the k algebra map f = evaluation at a; a k algebra map f:$k[X_1,...,X_n]$-->k is always surjective since it already is surjective on k, so f has a maximal ideal kernel M such that the composition k-->$k[X_1,...,X_n]$-->$k[X_1,...,X_n]/M \approx k$ is an isomorphism; a maximal ideal M such that the composition k-->$k[X_1,...,X_n]/M$ is an isomorphism is always of form $(X_1-a_1,...,X_n-a_n)$ where a_i = the unique element of k such that $X_i = a_i$, mod M; finally a maximal ideal of form M = $(X_1-a_1,...,X_n-a_n)$ determines the point a by setting a_i equal again to the unique element of k congruent mod M to X_i.

We want to prove the nullstellensatz which will simplify these correspondences as follows, in the case where k is algebraically closed:

Theorem: If k is algebraically closed the following are all true, and equivalent.

(i) Every proper ideal of $k[X_1,...,X_n]$ has a common zero in k^n.

(ii) For every maximal ideal M of $k[X_1,...,X_n]$, the composition $k\text{-->}k[X_1,...,X_n]/M$ is an isomorphism.

(iii) Every maximal ideal M of $k[X_1,...,X_n]$ has form $M = (X_1-a_1,...,X_n-a_n)$.

Cor: There is a one - one correspondence between points of k^n and maximal ideals of $k[X_1,...,X_n]$, with the point a corresponding to the kernel of evaluation at a.

I like Zariski`s version as follows which applies to all fields.

Zariski`s nullstellensatz:
If k is any field and M in $k[X_1,...,X_n]$ is any maximal ideal, the field $k[X_1,...,X_n]/M = E$ is a finite algebraic extension of k.

Cor: If k is algebraically closed, the inclusion k in $k[X_1,...,X_n]/M = E$ is an isomorphism, so $M = (X_1-a_1,...,X_n-a_n)$ for some a in k^n.

To prove Zariski`s version, all we need is basic facts about integral extensions, i.e. the ring theoretic analog of algebraic extensions of fields.

Proof of Hilbert's Nullstellensatz, (uncountable case only)

If k is a field, we define an affine algebraic set as a subset of k^n which is the common zero locus of some non empty collection of polynomials $\{f_i\}$ in $k[X_1,....,X_n]$. It is easy to see that the common zero locus of the collection $\{f_i\}$ is the same as the common zero locus of the ideal generated by the set $\{f_i\}$. We will denote this zero locus by $Z(\{f_i\})$ for zero locus, or $V(\{f_i\})$ for variety. The most basic question we can ask is whether the zero locus of a given set of polynomials is empty, i.e. whether a system of equations has any common solutions. If k is an arbitrary field, even Q, this can be a very difficult question. If k is algebraically closed there is a simple answer due to Hilbert, i.e. the common zero locus of a collection $\{f_i\}$ of polynomials in $k[X_1,....,X_n]$ is empty if and only if the ideal generated by the collection $\{f_i\}$ is the unit ideal.

If there is only one polynomial f in the collection, the proof is an easy induction. For n = 1 it is just the definition of an algebraically closed field. I.e. f has no zeroes if and only if f is a non zero constant, if and only if f is a unit, iff f generates the unit ideal. Note also that a non constant polynomial has zeroes but only a finite number of them, so since an algebraically closed field k is infinite, there are points of k where the polynomial vanishes and also points where it does not vanish. This is crucial for the induction. (Alternatively, note that if f were a non constant polynomial which vanished everywhere then 1+f would be non constant and have no zeroes, contradicting algebraic closure of the field.) Now by induction we may assume that a non constant polynomial in $k[X_1,....,X_{n-1}]$ must have zeroes but does not vanish everywhere, and write $f(X_1,....,X_n)$ as a polynomial in X_n, with coefficients which are polynomials in $k[X_1,....,X_{n-1}]$. If our polynomial f is "constant in X_n", i.e. if X_n does not occur in f, then we are done by induction, i.e. our polynomial vanishes at some but not all points. If X_n does occur in f, by induction there are points of k^{n-1} where the coefficient polynomial of the leading power of X_n does not vanish. Evaluating the coefficient polynomials at such a point we get a non constant polynomial in X_n with constant coefficients, hence there exist a non empty but only a finite set of values of X_n making the polynomial zero. Thus our original polynomial has zeroes, but does not vanish everywhere. **QED.**

Remark: Geometrically this inductive argument says that either the hypersurface $\{f=0\}$ of k^n is a cylinder over an algebraic set in k^{n-1}, or there are points of k^{n-1} over which the hypersurface has a finite non empty set of points. I.e. if we project the hypersurface into k^{n-1}, the image of the projection is non empty, and either the image is an algebraic subset of k^{n-1} or else there is a proper algebraic subset Z of k^{n-1} such that there are finitely many points of the hypersurface over each point of the complement of Z.

Thus projection of our hypersurface into k^{n-1} represents it either as a cylinder, or generically as a finite cover of some open subset. In general if we write f as a polynomial in X_n, with coefficients which are polynomials in $X_1,...,X_{n-1}$, there are infinitely many points of the hypersurface over each common zero in k^{n-1} of the coefficient polynomials.

If the coefficient polynomials of the non constant terms in X_n generate the unit ideal in $k[X_1,...,X_{n-1}]$, the projection is surjective, and there are are finitely many points over each point of k^{n-1}. We call the projection a "quasi finite" surjection in that case. If the coefficient of the highest power of X_n is a unit, e.g. if f is monic in X_n of degree d, then there are exactly d points (counted with multiplicities) over each point of k^{n-1}. In that case the projection is called a "finite" surjection.

For example, the projection of the hyperbola $\{XY-1 = 0\}$ onto the X axis is quasi finite (but not surjective), while the projection of the parabola $\{Y^2 - X = 0\}$ onto the X axis is finite, and surjective. The difference is the hyperbola has a vertical asymptote (over the point $X = 0$ where the coefficient of the leading power of Y vanishes) but the parabola does not.

For sets of more than one polynomial the result takes more work.

Theorem: ("weak nullstellensatz") The common zero locus of a (non empty) collection $\{f_i\}$ of polynomials in $k[X_1,....,X_n]$ is empty if and only if the ideal generated by the collection $\{f_i\}$ is the unit ideal. I.e. $V(\{f_i\})$ is empty if and only if 1 can be written as a linear combination $1 = ‡ g_if_i$ of the polynomials $\{f_i\}$, with coefficients g_i in $k[X_1,....,X_n]$.

Proof: There is an easy proof for uncountable, algebraically closed fields. Since this covers the complex numbers C, but not the algebraic closures of Q or Z_p, we will give it now and then give a more general proof using the theory of integral ring extensions in a later lecture on finite maps.

One direction is trivial: if 1 can be written as a linear combination $1 = \sum g_if_i$ of the polynomials $\{f_i\}$, then any common zero of the $\{f_i\}$ would be a zero of the polynomial 1, a contradiction. So assume the set $\{f_i\}$ does not generate the unit ideal. Then they all belong to some maximal ideal m. We will show that for any maximal ideal m in $k[X_1,....,X_n]$, there exist constants $a_1,....,a_n$ such that the polynomials $\{X_1-a_1,....,X_n-a_n\}$ all belong to m. Since the ideal generated by the X_i-a_i is maximal, being the kernel of the map $k[X_1,....,X_n]$-->k given by evaluation at a, this will prove that $m = (X_1-a_1,....,X_n-a_n)$. Then the point $a = (a_1,....,a_n)$ is a common zero of all polynomials in m, including therefore the $\{f_i\}$. (Note that any polynomial f in $k[X_1,....,X_n]$ can be written as $f(X) = \sum (X_i-a_i) g_i(X) + f(a)$. To prove this note it is true for $f = X_i$ and f = constant, and is also true for sums and products of polynomials for which it is true. Hence it is true for all polynomials. To obtain such an expansion just put Xi = Xi – ai +ai and expand.)

Now we prove that if m is any maximal ideal in $k[X_1,....,X_n]$, there exist elements $a_1,....,a_n$ such that for all i, X_i-a_i belongs to m. If n = 1, then $k[X_1]$ is a pid, so m is prime and principal, hence is generated by one irreducible polynomial. Since k is algebraically closed an irreducible polynomial must be linear, and after multiplication by a unit, can be taken to be monic, hence of form X_1-a_1. That does the case n = 1.

Now let m be maximal in $k[X_1,....,X_n]$ with n > 1, and for each i consider the map $k[X_i]$-->$k[X_1,....,X_n]/m$. Since m is maximal hence prime, the kernel of this map is a prime ideal of $k[X_i]$. If the kernel is non zero, then it is a maximal ideal of $k[X_i]$, hence contains an element of form X_i-a_i which then lies in m, as desired. So we must rule out the possibility that the kernel of the map could be zero. If the kernel of the map $k[X_i]$-->$k[X_1,....,X_n]/m$ were zero, then $k[X_i]$ would be embedded as a subring in the field $k[X_1,....,X_n]/m$, and hence its field of fractions $k(X_i)$ would be embedded as a subfield and hence as a k vector subspace. Since the target field $k[X_1,....,X_n]/m$ is a finitely generated algebra over k, it is a countable dimensional vector space over k. We will show that $k(X_i)$ cannot be embedded in $k[X_1,....,X_n]/m$ if k is uncountable, by showing that $k(X_i)$ is an uncountable dimensional vector space over k.

To do this we will prove the uncountable subset $\{1/(X_i-c)$: all c in k$\}$ is independent over k. For simplicity we write X for X_i. If $\sum a_j/(X-c_j) = 0$ is a linear dependency relation in k(X) where all c_j are distinct, then multiplying by $\prod_j (X-c_j)$ gives $\sum_j (a_j \prod_{t\neq j}(X-c_t)) = 0$. If we set $X = c_j$, all terms of this sum except the jth one vanish, so we get $a_j (\prod_{t\neq j} (c_j-c_t)) = 0$. Since for $t \neq j$, $(c_j-c_t)) \neq 0$, we must have $a_j = 0$. Thus the original linear dependency relation was the trivial one. This proves the nullstellensatz for uncountable fields. **QED.**

This result tells us that over an algebraically closed field k, every ideal of $k[X_1,....,X_n] = k[X]$ except the unit ideal (1), defines a non empty algebraic subset of k^n. Thus there is a surjection from the collection of non trivial or “proper” ideals of k[X], to the collection of non empty algebraic subsets of k^n, taking J to V(J). To know a bit more about this parametrization of algebraic sets by ideals, we should also ask when two different ideals define the same subset. Notice that an algebraic set V is defined by a unique largest ideal, the ideal I(V) of all polynomials vanishing on the set V. I.e. if J is any ideal, and V(J) the associated algebraic set, then I(V(J)) is the largest ideal defining the set V(J). In particular, V(I(V(J))) = V(J), and I(V(J)) contains J. There is no reason for I(V(J)) to equal J. Indeed if m = $(X_1-a_1,....,X_n-a_n)$, then the point a = $(a_1,....,a_n)$ is defined by any one of the infinite sequence of ideals m, m^2, m^3, m^4,........, of which m is the unique largest one. Note that if J is any ideal and f is in the "radical" of J, i.e. f^r belongs to J for some $r \geq 1$, then f vanishes on every point of V(J). Thus I(V(J)) contains the radical of J. The "strong nullstellensatz" of Hilbert says that the converse is true over an algebraically closed field k, and hence that algebraic subsets of k^n corresponds 1-1 with radical ideals of k[X], i.e. those ideals which are equal to their own radical. By a well known trick, this result is a corollary of the weak nullsatz. First we deduce the result using the trick and then explain how the trick may be motivated geometrically.

Theorem (strong nullstellensatz)**:** If k is an algebraically closed field, i.e. a field for which the weak nullstellensatz is true, and if J is any ideal of $k[X_1,....,X_n]$, then a polynomial f in k[X] vanishes on the set of common zeroes of J if and only if some positive power of f belongs to J.
Proof: If J = (f_1,..,..,f_r) in $k[X_1,....,X_n]$, and f vanishes on V(J), we add one more variable X_{n+1} and consider the ideal $(f_1,....,f_r, X_{n+1}f - 1)$ in $k[X_1,....,X_{n+1}]$. By hypothesis, the last generator of this ideal equals -1 on the set V(J) where the first r generators vanish, hence the new ideal has no common zeroes. Thus by the nullstellensatz for the ring $k[X_1,....,X_{n+1}]$, the new ideal is the unit ideal. Thus there exist polynomials $g_1,....,g_r$, g in $k[X_1,....,X_{n+1}]$ such that $\sum f_ig_i + (X_{n+1}f-1)g = 1$. Now map the polynomial ring $k[X_1,....,X_{n+1}]$ into the fraction field of $k[X_1,....,X_n]$ sending X_{n+1} to 1/f. (If f = 0, there is nothing to prove, since then f belongs to J.) Under this k algebra map, 1 goes to 1, and $(X_{n+1}f-1)g$ goes to 0, so the equation $\sum f_ig_i + (X_{n+1}f-1)g = 1$ in $k[X_1,...,X_{n+1}]$ yields an equation $\sum f_i(X_1,...,X_n)g_i(X_1,...,X_n, 1/f) = 1$ in $k(X_1,...,X_n)$, where the only denominators are powers of f. If we multiply by a suitably high power of f, say f^s, this eliminates all denominators, and yields an equation $\sum f_i(X_1,...,X_n)h_i(X_1,...,X_n) = f^s$, as desired. **QED.**

Corollary: If k is algebraically closed, then assigning an ideal J to its set of common zeroes V(J), and assigning an algebraic set V to its corresponding ideal I(V), are mutually inverse (1-1) correspondences between the collection of all radical ideals of $k[X_1,....,X_n]$ and the collection of all algebraic subsets of k^n.

Exercise: The collection of all algebraic subsets of k^n is closed under taking arbitrary intersections and finite unions, hence forms the family of closed sets for a topology on k^n called the Zariski topology. The induced topology on a (closed) algebraic set V is likewise called the Zariski topology for V.

We can refine our correspondences further to determine which algebraic sets correspond to prime ideals. Note in particular that every prime ideal is radical.

Definition: An algebraic set Y is called "irreducible" if and only if whenever V,W are algebraic sets such that Y = V U W, (V union W), then either Y = V or Y = W. Equivalently, whenever Y is contained in a union V U W, of algebraic sets then Y is contained in either V or W.

Proposition: An algebraic set Y in k^n is irreducible if and only if its ideal I(Y) in $k[X_1,....,X_n]$ is prime.
proof: If I(Y) is not prime there exist f,g such that fg is in I(Y) but neither f nor g vanishes on Y. Then Y is contained in V(f) U V(g), but not in either V(f) or V(g), so Y is reducible. Conversely if Y is reducible, then Y = V U W where both V,W are strictly smaller than Y. Thus I(Y) is the intersection of I(V) and I(W), ideals which are strictly larger than I(Y). Then if f belongs to I(V) but not to I(Y), and if g belongs to I(W) but not to I(Y), then fg belongs to the intersection of I(V) and I(W), i.e. to I(Y), which therefore is not prime. **QED.**

Thus when k is algebraically closed, we have 1-1 correspondences as follows:

{maximal ideals of k[X]} <--------> {points of k^n}

{prime ideals of k[X]} <--------> {irreducible algebraic sets of k^n}

{radical ideals of k[X]} <----------> {all algebraic subsets of k^n}

Notice these correspondences are inclusion reversing, i.e. a larger ideal has a smaller set of common zeroes, and a larger algebraic set has a smaller ideal of functions which vanish on it.

We obtain as corollaries information about the coordinate rings of affine algebraic sets. We often use the phrase "algebraic variety" or just "variety" to describe an algebraic set, but sometimes these terms are used to refer only to an irreducible algebraic set, so be careful.

Definition: If Y is an affine algebraic subset of k^n with ideal I(Y), the “coordinate ring of Y” is defined to be $k[Y] = k[X_1,...,X_n]/I(Y)$ = the k algebra of functions on Y which are restrictions of polynomial functions from $k[X_1,...,X_n]$.

Remark: It follows from the weak nullstellensatz that if Y is an affine algebraic set in k^n, for an algebraically closed field k, there is a 1-1 correspondence between points of Y and maximal ideals of $k[X_1,...,X_n]$ that contain I(Y), i.e. with all maximal ideals of k[Y].

Definition: A polynomial map between two algebraic sets Y, Z in k^n, k^m respectively, is defined by m polynomial functions on Y that map Y into Z.

Useful Exercises:

Given a polynomial map p from Y to Z, we get a k algebra map from k[Z] to k[Y] by pulling back functions by composing with p. Conversely, every such k algebra map k[Z] --> k[Y], is induced by a unique polynomial map Y-->Z.

(This is true for any field, algebraically closed or not. One way to check this is by using the image of the coordinate functions k[Z] --> k[Y] to define the map Y-->Z. A more intrinsic way is as follows. When k is not algebraically closed, not every maximal ideal of k[Y] arises from a point of Y. Those maximal ideals that do arise from points are exactly the kernels of k algebra homomorphisms k[Y]-->k, or equivalently those maximal ideals m such that the composition k-->k[Y]/m is an isomorphism. Then given a k algebra homomorphism k[Z]-->k[Y], we can define a map back on points Y-->Z, by pulling back maximal ideals of this special type, e.g. by composing k algebra homomorphisms k[Z]-->k[Y]-->k. In particular, two affine algebraic sets Y,Z over any field k, are isomorphic (via mutually inverse polynomial maps) if and only if their coordinate rings k[Z], and k[Y] are isomorphic as k algebras. (Is this correct? Please check me.))

Corollary: (i) The ring k[V] of an affine algebraic set V is a finitely generated k algebra without nilpotents, and if k is algebraically closed every such k algebra occurs as k[V] for some V.
(ii) The ring k[V] of an affine algebraic set V is a domain if and only if V is irreducible.
proof: In (i) to show every fin. gen. k algebra R without nilpotents occurs as k[V], just map some $k[X_1,...,X_n]$ onto R with kernel J, and let V = V(J). **QED.**

Thus if k is algebraically closed, the category of irreducible affine algebraic sets over k, is equivalent to the category of finitely generated k algebras which are domains, but under this equivalence, morphisms in the two categories go in opposite directions.

These results allow a definition of an abstract affine algebraic set, independent of a a particular embedding.

Definition: An (abstract) affine variety over an algebraically closed field k, is a pair (Y,R), consisting of a set Y and a finitely generated k algebra R of k valued functions on Y, such that the map taking a point p of Y to the maximal ideal of functions in R vanishing at p, is a 1-1 correspondence between points of Y and maximal ideals of R.

Example: Let Y be an algebraic subset of k^n (irreducible for simplicity), k algebraically closed, and f an arbitrary non zero function in k[Y]. Then we claim the "principal" open subset $U_f = \{f \neq 0\}$ of Y, is an abstract affine variety, with coordinate ring $k[U_f] = k[Y][1/f]$, thought of as a subring of the fraction field of k[Y]. I.e. this is a finitely generated k algebra of functions on U_f, and its maximal ideals correspond exactly to those maximal ideals of k[Y] which do not contain f, i.e. to points of U_f.

To obtain an embedding of U_f as a closed subset of some affine space, take the generators $g_1,...,g_n$ of k[Y] corresponding to the variables $X1,...,X_n$, and add in the additional generator 1/f to get generators for $k[U_f]$. Then mapping the polynomial ring $k[X_1,...,X_n,X_{n+1}]$ onto $k[U_f]$ by sending X_1 to $g_1,...,X_n$ to g_n, and X_{n+1} to 1/f, defines a k algebra map U_f-->k^{n+1}, taking $(x_1,...,x_n)$ to $(x_1,...,x_n, 1/f(x_1,...,x_n))$. The map is injective and the image is the closed subset of points $(x_1,...,x_n,x_{n+1})$ in k^{n+1} such that $x_{n+1}.f(x_1,...,x_n) -1 = 0$. Thus this image set is an affine closed subset V of k^{n+1}. Since the map back on rings of functions takes X_1 to $g_1,...,X_n$ to g_n, and X_{n+1} to 1/f, it maps the ring $k[X_1,...,X_{n+1}]/(X_{n+1}.f(X_1,...,X_n) -1)$ of V isomorphically to $k[U_f]$. Thus as maps of abstract affine varieties, this is an isomorhism.

This explains the trick of Rabinowitz in the proof above that the weak nullstellensatz implies the strong nullstellensatz. I.e. the fact that f vanishes on the zero locus of the f_i in that proof, means that the f_i have no zeroes on the set U_f where f does not vanish. Translating this statement to the isomorphic affine closed set V in one dimension higher, we have that a certain collection of functions have no common zeroes in k^{n+1}. Then the weak nullsatz says those functions generate the unit ideal there. This then yields the original claim one dimension lower.