THE NULLSTELLENSATZ; CLOSED POINTS AND k-POINTS

PETE L. CLARK

1. The Nullstellensatz

Our treatment of affine varieties up to this point has been "nonstandard" in that we have refused to associate to an affine variety a set of points. Recall that the standard definition of an affine variety (or rather, an affine subvariety of $\mathbb{A}^n$) over an algebraically closed field k is as a subset of k^n which is defined as the set $Z(S)$ of simultaneous zeros of a family $S = \{P_i\} \subset R_n := k[t_1, \ldots, t_n]$ of polynomials. If $I(S)$ is the ideal generated by S, one sees immediately that $Z(S) = Z(I(S))$. Moreover, by the Hilbert Basis Theorem, every ideal of R_n is finitely generated, so even though *a priori* we allow intersections of possibly infinite sets of polynomials, any such set can also be "cut out" by a finite set of polynomials.

We can also go in the other direction: given an ideal I in R_n, we can consider the set $V(I)$ of all points of $x \in \mathbb{A}^n$ such that $f \in I \implies f(x) = 0$. For instance, associated to the ideal $(xy) \in \mathbb{C}[x, y]$, we get the set of all points for which $x = 0$ or $y = 0$, i.e., two lines meeting at the origin.

Thus we have two sets: the set of all subsets of $\mathbb{A}^n$, and the set of all ideals of R_n, and we have two mappings between them, V and I. The following formal properties satisfied by V and I are immediate:

(GC1) $S_1 \subset S_2 \implies I(S_2) \subset I(S_1)$, $I_1 \subset I_2 \implies V(I_2) \subset V(I_1)$.

That is, both maps are **order reversing** (or "antitone").

(GC2) For $S \subset \mathbb{A}^n$ and I an ideal of R_n we have:

$$S \subset V(I) \iff I \subset I(S).$$

Indeed, both conditions say: $\forall f \in I,\ \forall P \in S,\ f(P) = 0$.

In general, if we have two partially ordered sets S and T and maps $f : S \to T$, $g : T \to S$ satisfying (GC1) and (GC2) (with "$\subset$" replaced by the "$\leq$" in the posets), we say that (S, T, f, g) form a **Galois connection**.[1]

Notice that the concept of an antitone Galois connection is inherently symmetrical in the sense that if (S, T, f, g) is a Galois connection, so too is (T, S, g, f). (An "isotone" Galois connection breaks this symmetry, and represents the fundamental asymmetry between left and right adjoint functors.)

[1]More precisely, an "antitone" Galois connection. When f and g are both order-preserving, an appropriate modification of (GC2) leads to the notion of an "isotone" Galois connection.

Proposition 1. *Let (S, T, f, g) be a Galois connection on partially ordered sets.*
a) For $s \in S$ and $t \in T$, $f(g(f(s)) = s$, $g(f(g(t)) = t$.
*b) The operators $c_S : S \to S$, $c_T : T \to T$ by $c_S : s \mapsto f(g(s))$, $c_T : t \mapsto g(f(s))$ are both **closure operators** in the sense of order theory.*
c) Let $S_c = g(f(S)) = \{s \mid c_S(s) = s\}$, $T_c = f(g(T)) = \{t \mid c_T(t) = t\}$ be the subsets of "closed" elements. Then f and g induce mutually inverse order reversing bijections between S_c and T_c.

Exercise: Prove Proposition X.X. (The hardest part will probably be to figure out what a "closure operator in the sense of order theory" is supposed to mean. Hint: wikipedia has the correct definition.)

This little bit of structure pulled from thin air will quicken the heart of any Bourbakiste. But beyond the formalism, the key question is: exactly which sets are closed? Without knowing this, we haven't proved the Nullstellensatz any more than the analogous formalities between sets and groups of automorphisms prove the Galois correspondence for Galois field extensions.

Here are two general observations we can make without assuming that the ground field is algebraically closed:

Every "closed" ideal – i.e., an ideal of the form $I(S)$ – is a radical ideal.

Indeed, an ideal I is radical if $f^n \in I$ implies $f \in I$. But if f^n vanishes identically on S, then so does f.

The closed subsets of $\mathbb{A}^n$ are closed under arbitrary intersections (including the "empty intersection": $\mathbb{A}^n = V((0))$ and under finite unions (including the "empty union": $\emptyset = V(\{1\}) = V(R_n)$, and therefore form the closed sets for a unique topology on $\mathbb{A}^n$, the **Zariski topology**.

Exercise: a) Prove these facts.
b) Prove or disprove: the Zariski topology on $\mathbb{A}^n_{/k}$ coincides with the topology it inherits as a subset of $\mathbb{A}^n_{/\overline{k}}$.
c) Show that the Zariski topology is T_1: i.e., singleton subsets are closed.
d) Show that when $n = 1$, the Zariski topology is the coarsest T_1 topology on k: namely, the topology in which a proper subset is closed iff it is finite.
e) For any $n \geq 1$, show that the Zariski topology on k^n is discrete iff k is finite.
f) For any infinite field and $m, n \geq 1$, show that the Zariski topology on k^{m+n} is strictly finer than the product of the Zariski topologies on k^m and k^n.

Remark: It is often lamented that the Zariski topology (especially when $k = \mathbb{C}$) is distressingly "coarse". It is true that it is much coarser than the "analytic topology" on k^n when k is a topological field (i.e., the product topology from the topology on k). But in fact from an algebraic perspective the Zariski topology is if anything too fine: we will see why later on when we extend the topology to all prime ideals on an affine algebra. But also the fact that the Zariski topology on a $\mathbb{F}_q^n$ is discrete creates many geometric problems.

Exercise: Let k be a field, $n \in \mathbb{Z}^+$ as above. Explicitly compute the ideal $I(k^n)$, i.e., the set of all polynomials which vanish at *every* point of k^n. Do we necessarily have $I(k^n) = \{0\}$? (Hint: No, but we can figure out exactly for which fields this occurs.)

Let us now state the "classical case" of the Nullstellensatz:

Theorem 2. *(Classical "Strong" Nullstellensatz) Let k be an algebraically closed field, let $R_n = k[x_1, \ldots, x_n]$. Then:*
a) For any Zariski-closed subset $S \subset k^n$, $V(I(S)) = S$.
b) For any ideal J of R_n, $I(V(J)) = \mathrm{rad}(J)$.
*Thus there is an inclusion-reversing, bijective correspondence between Zariski-closed subsets of k^n and **radical** ideals of R_n.*
c) I induces a bijective correspondence between the singleton sets of k^n and the maximal ideals. Explicitly,

$$V((a_1, \ldots, a_n)) = \langle x_1 - a_1, \ldots, x_n - a_n \rangle.$$

Exercise: Deduce part c) directly from parts a) and b). Note that the Nullstellensatz comes very close to saying that (when $k = \overline{k}$) there is no loss of information in identifying a "presented affine k-algebra" $A = l[x_1, \ldots, x_n]/I$ with the corresponding subset $V(I) \subset k^n$. This is certainly the more traditional approach, and the affine algebra $R_n/I(S)$ is traditionally called the **coordinate ring** of the affine variety S and denoted $k[S]$. Here we say "comes very close" because we only get reduced affine algebras in this way. At the moment we haven't given much in the way of concrete justification for considering affine algebras with nilpotent elements, so it's not yet so clear whether this is a drawback or a "feature" of the theory.
Note also that it makes excellent sense to view $k[S]$ as the ring of all functions on S. Here, we think of the functions on S as the ring of all functions $S \to k$ which are the restrictiion to S of at least one polynomial $f \in R_n$. Note the "at least one": if $f, g \in R_n$ are two polynomials such that $(f - g)$ vanishes at every point of S, then $f - g$ is the zero function *on* S and thus $f|_S = g|_S$. Thus this ring of functions is precisely $R_n/I(S)$. We can define a regular function from $k^n \to k^m$ to be given by polynomials $(P_1, \ldots, P_m)$, and then if $V \subset k^n$ and $W \subset k^m$, we can define regular functions from V to W.

Exercise: Do so. Also show that the set of all functions from V to W is well-defined independent of the affine embeddings of V and W.

Now, what about the case of an arbitrary k? As we have seen, the Nullstellensatz cannot possibly hold in the above form, because the zero set of a proper ideal may be empty. Indeed, if k is not algebraically closed, then by definition there exists a nonconstant polynomial $P(x_1) \in k[x_1]$ without any roots, and then $V(\langle P(x_1) \rangle) = \emptyset$.

So what to do? It turns out the result we want is the following:

Theorem 3. *(Modern "Weak" Nullstellensatz) Let k be an arbitrary field, A a finitely generated k-algebra, and $\mathfrak{m}$ a maximal ideal of A. Then $A/\mathfrak{m}$ is a finite-dimensional k-vector space.*

This was probably not what you're expecting, and a natural reaction is: what does this have to do with the previous Nullstellensatz?!? Let us show how to recover

Hilbert's Nullstellensatz from Theorem XX when $k = \overline{k}$. First, the assumption that k is algebraically closed is equivalent to: the only finite-degree field extension of k is k itself, so every maximal ideal $\mathfrak{m}$ of A has residue field k. To be precise, let $q : A \to A/\mathfrak{m} \xrightarrow{\sim} k$ be the composite map. Now put, for $i = 1, \ldots, n$, $a_i = q(x_i)$. Then $\mathfrak{m}$ contains $x_i - a_i$ for all i, and since the ideal generated by these n linear polynomials is evidently a maximal ideal, we must have $\mathfrak{m} = \langle x_1 - a_1, \ldots, x_n - a_n \rangle$.

Remark: In the classical case – $k = \overline{k}$, of course – the statement that all maximal ideals in R_n are of the form $\langle x_1 - a_1, \ldots, x_n - a_n \rangle$ has also been called the **weak Nullstellensatz**, because it is not straightforward to deduce the full "radical ideals are closed" version of the Nullstellensatz from this one. It is also traditional to call Theorem 3 a "weak Nullstellensatz", but this usage is more puzzling, since almost every text on the subject explains how the strong Nullstellensatz can be deduced from this result by a simple argument (**Rabinowitsch trick**). Let us now give a slightly modernized version of this deduction, to establish the following

Theorem 4. *(Hilbert-Jacobson Nullstellensatz) Let A be a finitely generated k-algebra. Then every prime ideal of A is the intersection of all maximal ideals containing A.*

Proof: Let $\mathfrak{p}$ be a prime ideal of A. It suffices to find, for each $a \in A \setminus \mathfrak{p}$, a maximal ideal $\mathfrak{m}$ containing $\mathfrak{p}$ and not containing a. Let $A_a := A[\frac{1}{a}]$, and let $\mathfrak{p}_a := \mathfrak{p}A_a$. Recall that $\mathfrak{p}_a$ remains a prime ideal in the localization since $\mathfrak{p}$ does not meet the multiplicative subset $\{a^n\}$. Oberserve also that A_a remains an affine k-algebra, being generated by a single element over A. Let $\mathfrak{m}_a$ be any maximal ideal of A_a containing $\mathfrak{p}_a$, and put $\mathfrak{m} := \mathfrak{m}_a \cap A$. Then by passing to the quotient we get an injection

$$A/\mathfrak{m} \to A_a/\mathfrak{m}_a.$$

Thus the integral k-algebra $A/\mathfrak{m}$ is embedded in the finite-dimensional k-algebra $A_a/\mathfrak{m}_a$. By the "Weak" Nullstellensatz, $\dim_k A_a/\mathfrak{m}_a$ is finite, hence so also is $\dim_k A/\mathfrak{m}$. This implies that $A/\mathfrak{m}$ is a field and hence $\mathfrak{m}$ is a maximal ideal – let us spell this out for completeness. Take $0 \neq x \in A/\mathfrak{m}$; by finite dimensionality, we get a nontrivial k-linear dependence relation between the powers of x, say

$$x^n + c_{n-1}x^{n-1} + \ldots + c_1 x + c_0 = 0, \ c_i \in k.$$

If we choose such a relation of minimal degree n, then we must have $c_0 \neq 0$, since otherwise we could factor out $x \neq 0$ and get a relation of smaller degree. Thus

$$x(x^{n-1} + c_{n-1}x^{n-2} + \ldots + c_1) = \frac{-1}{c_0},$$

and it is now clear that x is invertible.

Exercise: When $k = \overline{k}$, deduce Theorem 2 from Theorem 3 and Theorem 4. (Hint: recall that **in any ring** every radical ideal is the intersection of all prime ideals containing it.)

A commutative ring is called a **Hilbert-Jacobson ring**[2] if every radical ideal is the intersection of all maximal ideals containing it. Thus Theorem 4 precisely asserts that an affine algebra is a Hilbert-Jacobson ring, and this result has all the

[2]Or just a Hilbert ring, or just a Jacobson ring: all these terms are synonymous.

strength of the classical Nullstellensatz.

We remark that there are parts of mathematics other than algebraic geometry in which Hilbert-Jacobson rings occur naturally. For instance there are similar results for rings of analytic functions of various kinds. On the other hand, from a purely algebraic perspective, the Hilbert-Jacobson condition is very appealing and has been much studied. For instance, one generalized Nullstellensatz asserts that if R is a Hilbert-Jacobson ring, then so is any finitely generated R-algebra. Further implications are explored at length in Kaplansky's text on commutative rings.

1.1. **Proofs of the Modern "Weak" Nullstellensatz.**

The standard argument goes by way of an intermediate result of Emmy Noether. Since Noether's theorem has other applications in algebraic geoemetry, this is still a very popular and reasonable way to proceed. We give a version of Noether's theorem which is slightly stronger than the usual one:

Theorem 5. *(Noether normalization theorem) Let k be a field, and let A be an integral affine k-algebra, with fraction field K, of dimension d.*
a) There exist elements $x_1, \ldots, x_d \in A$ such that:
(i) The set $\{x_1, \ldots, x_d\}$ is algebraically independent over k; i.e., $k[x_1, \ldots, x_d]$ is a polynomial ring.
(ii) The inclusion $k[x_1, \ldots, x_d] \hookrightarrow A$ makes A into a finitely generated k-module.
b) Assume moreover that either k is perfect or A is geometrically integral. Then we can choose $x_1, \ldots, x_d \in A$ satisfying (i) and (ii) above and also:
(iii) The finite field extension $K/k(x_1, \ldots, x_d)$ is separable.

The geometric interpretation of this is that every integral affine variety of dimension d is a finite branched covering of affine d-space $\mathbb{A}^d$. Moreover, if k is geometrically integral (and, quite trivially, if k is perfect), then we can choose the branched covering to be separable.

A proof of part a) of Theorem 5 can be found in Liu's book: Proposition 2.1.9 on page 29. The deduction of all of the Nullstellensatzen (and especially Theorem 3) follows on pages 30-32. Note that the stronger part b) is not needed here. A proof of part b) can be found in Eisenbud's *Commutative Algebra* text.

Part b) will certainly be useful for us later though: stay tuned.

There has also been an industry of giving completely algebraic (and maximally elementary) proofs of the Nullstellensatz, which in particular avoid Noether's normalization lemma. To my mind, the slickest of the slick is to be found in Richard G. Swan's note *On Munshi's Proof of the Nullstellensatz* (a link is available on the course webpage), which disposes of the whole thing using a very simple and beautiful induction argument.

If you are working with specific algebraic varieties, it is obviously very desirable to be able to implement the Nullstellensatz algorithmically. None of the classical proofs are "effective" in this sense, and the most straightforward adaptations give very poor bounds on degrees of polynomials. Implementing an efficient algorithmic

Nullstellensatz is a major issue in computational algebraic geometry (unfortunately I have already told you all I know about it).

1.2. **The most general (and trivial) Nullstellensatz.**

What should the Nullstellensatz say for a general commutative ring R? In this case, it is not necessarily true that any radical ideal (or even any prime ideal) is the intersection of the maximal ideals containing it. Consider for example the case of a DVR where there is the zero ideal and a unique maximal ideal.

However, there is again a Galois connection here. Let $\mathcal{I}(R)$ be the set of all ideals of R, and let $\operatorname{Spec} R$ be the set of all prime ideals of R. Then to an ideal I we associate $V(I)$, the set of all prime ideals containing I; whereas to any set S of prime ideals of R, we associate $I(S) = \bigcap_{\mathfrak{p} \in S} \mathfrak{p}$. The following facts are all immediate:

Theorem 6. *Let R be any commutative ring. The pair (V, I) forms a Galois connection between the set of ideals of R and the power set of the set of the set of all prime ideals of R. The closed ideals are precisely the radical ideals, and the closure $S \mapsto \overline{S} = V(I(S))$ is the closure operator for a unique topology on the set of prime ideals, called the Zariski topology. Thus we get a bijective correspondence between radical ideals and Zariski-closed subsets of* $\operatorname{Spec} R$.

Note that the algebraic content here is hardly more than the statement that the radical of any ideal is equal to the intersection of all the primes containing it. What right do we have to regard this as a Nullstellensatz?

The first point is that we get to consider $\operatorname{Spec} R$ as a topological space. In particular, a prime ideal $\mathfrak{p}$ of R is regarded as a "point" in this new space. The term spectrum is taken from the Gelfand spectrum in analysis: if X is a compact Hausdorff space, let $R_X = \mathbb{C}(X, \mathbb{C})$ be the ring of continuous, complex-valued functions on X. For every point $x \in X$, the set $\mathfrak{m}_x := \{f \in R_X \mid f(x) = 0\}$ is a maximal ideal of R_X. Conversely, it can be shown that every maximal ideal is of this form. Thus the set of maximal ideals of X is in bijection to X (and there is a way of giving the set a topology which makes this bijection a homeomorphism; for lack of time we omit this), so overall the ring R_X is recovered as the ring of all continuous, $\mathbb{C}$-valued functions on its (maximal) spectrum $\operatorname{MaxSpec}(R_X)$.

We would like to do something similar with R and $\operatorname{Spec} R$ (and for that matter, with R and $\operatorname{MaxSpec} R$ when R is a Jacobson ring). Let $f \in R$; how do we view f as a "function" on $\operatorname{Spec} R$? The idea here is that given a point $\mathfrak{p} \in \operatorname{Spec} R$, we can form the homomorphism $R \to R/\mathfrak{p}$ to an integral domain, and then take the field of fractions, overall getting a homomorhpism $R \to k(\mathfrak{p})$, where $k(\mathfrak{p})$ is the fraction field of $R/\mathfrak{p}$, also called the **residue field** at $\mathfrak{p}$. In case R is an affine algebra over an algebraically closed field k and $\mathfrak{p}$ is a maximal ideal, then we know that $k(\mathfrak{p}) = k$ (Nullstellensatz!). Thus it is literally true that f is a function from $\operatorname{MaxSpec}(R)$ to $\mathbb{C}$. Of course this picture is compatible with f being a regular function on the corresponding affine variety.

In the general case we need to loosen our notion of a "function" a bit, so as to allow the codomain to vary from point to point. In other words, we regard the

associated function to $f \in R$ to be the function $\mathfrak{p} \in \operatorname{Spec} R \mapsto \overline{f} \in kk(\mathfrak{p})$. In this way **every** commutative ring can be viewed as a ring of functions (of a sort) on a topological space, its prime spectrum $\operatorname{Spec} R$.

Thus we can enlarge our category of affine k-varieties to a category of affine schemes: it is just the opposite of the category of commutative rings, where the morphisms are ring homomorphisms. To be precise, we can also, for any commutative ring k, define the category of affine k-schemes, which is opposite to the category of commutative k-algebras and k-algebra homomorphisms. In this view we recover the category of affine schemes as that of affine $\mathbb{Z}$-schemes, since every commutative ring has the canonical structure of a $\mathbb{Z}$-algebra.

But we have also done a little bit more: to an affine scheme we have associated a set of points $\operatorname{Spec} R$, which is also endowed with the structure of a topological space via the Zariski topology. This is our first notion of a point on a scheme.

The Zariski topology has several important properties, some nice and some which appear at first to be somewhat strange. First $\operatorname{Spec} R$ is always **quasi-compact**, i.e., every open covering admits a finite subcovering. Moreover, if R is Noetherian, then $\operatorname{Spec} R$ has the alarmingly strong property of being a **Noetherian topological space**.

Proposition 7. *For a topological space X, TFAE:*
(i) X satisfies the ascending chain condition (ACC) on open subsets.
(ii) X satisfies the descending chain condition (DCC) on closed subsets.
(iii) Every open subset of X is quasi-compact.
(iv) Every subset of X is quasi-compact.
*A space satisfying any – hence all – of these properties is called **Noetherian**.*

Since in a Hausdorff space quasi-compact subsets are always closed, and a space in which every subspace is closed is discrete, we see that a Noetherian Hausdorff space is both quasi-compact and discrete, hence finite. Thus for any Noetherian ring with infinitely many prime ideals, $\operatorname{Spec} R$ cannot be a Hausdorff space. It satisfies some weaker properties: it is a **Kolmogorov space** (or T_0), and every irreducible closed subsets is the closure of a unique point. A space with these two properties is called **sober**.

Let R be an arbitrary commutative ring, and let $\mathfrak{p} \in \operatorname{Spec} R$. Then the closure of the singleton set $\{\mathfrak{p}\}$ is the set of all prime ideals containing $\mathfrak{p}$. It follows that $\{\mathfrak{p}\}$ is closed iff $\mathfrak{p}$ is a maximal ideal. Thus maximal ideals correspond to closed points, and in particular $\operatorname{Spec} R$ is a T_1-space (all points are closed) iff R is zero-dimensional.[3]

Thus, although it does not play as crucial a role in the case of arbitrary rings, we may as well also endow $\operatorname{MaxSpec} R$ with the Zariski topology – the topology obtained by the Galois connection restricted to sets of maximal ideals is the same as the subspace topology from $\operatorname{Spec} R$. ten $\operatorname{MaxSpec} R$ is a T_1-subspace but is no

[3]It can be shown that when R is zero-dimensional, $\operatorname{Spec} R$ is in fact Hausdorff, but this is a bit tricky and not directly relevant to our geometric ambitions.

longer sober: for example, if k is any field and A is any regular, integral affine curve – e.g. $k[t]$ – then $\operatorname{MaxSpec} k[t]$ is endowed with the cofinite topology. If there are infinitely many maximal ideals – e.g. when k is infinite and $A = k[t]$, or when $k = \overline{k}$ – then the cofinite topology is T_1, not Hausdorff, and has no generic point. In $\operatorname{Spec} A$, in contrast, there is also the generic point η corresponding to the zero ideal, whose closure is the entire space, and for which the open neighborhoods are those with finite complement. In particular η is itself neither open nor closed, and thus $\operatorname{MaxSpec}(A)$ is not closed in $\operatorname{Spec} A$.

Remark: A natural and interesting question is to characterize the topological spaces which occur (up to homeomorphism, of course) as $\operatorname{Spec} R$ for some commutative ring R, or, if you like, the **spectral spaces**. This was in fact done in the Princeton thesis[4] of Mel Hochster.

2. More on Closed Points on Affine Varieties

Let A be an affine k-algebra. We have given a meaning to the set of closed points of A: it is the topological space $\operatorname{MaxSpec}(A)$, and we have given a Galois connection between subsets of $\operatorname{MaxSpec}(A)$ and ideals of A which is a reasonable generalization of the classical Nullstellensatz.

Nevertheless, we are still missing some of the "extrinsic geometry" that exists in the classical case (i.e., $k = \overline{k}$) in which via a choice of "presentation" (i.e., surjective homomorphism) $k[t_1, \ldots, t_n] \to k[t]/J = A$ we recover $\operatorname{MaxSpec}(A)$ as the subset $V(J)$ of k^n. To put it plainly, we can see the maximal ideals of A just as n-tuples of points.

A moment's thought reveals that in case of an arbitrary k, the set $V(J)$ corresponds bijectively to maximal ideals $\mathfrak{m}$ of A with residue field k; whereas, as we have noted already, whenever $k \neq \overline{k}$ there will always be maximal ideals whose residue field is a proper finite extension l of k. How do we "see" these points in affine space?

Clearly some sort of base change is called for, but in order to get a precise result we need to set things up carefully. On the level of sets, the following construction is appealing: let $G = \operatorname{Aut}(\overline{k}/k)$. If k is perfect, this is called the **absolute Galois group** of k; whereas if k is not perfect, the absolute Galois group is by definition $\operatorname{Aut}(k^{\mathrm{sep}}/k)$, but – since every element of $\operatorname{Aut}(k^{\mathrm{sep}}/k)$ extends uniquely to an automorphism of $\overline{k}$ – these groups are in fact isomorphic. Now G has a natural (i.e., coordinatewise) action on $\overline{k}^n$. If J is any ideal of $k[t_1, \ldots, t_n]$, then $\overline{J} = J\overline{k}[t_1, \ldots, t_n]$ is an ideal of $\overline{k}[t_1, \ldots, t_n]$ defined by polynomials with k-coefficients, so its zero set $V(\overline{J})$ is G-invariant.

If we assume that k is perfect, then we have the appealing identity

$$V(J) = V(\overline{J})^G,$$

[4]Written under the direction of Goro Shimura. For more than 50 years Shimura has been one of the most powerful and influential of mathematicians (full stop!). However, to the best of my knowledge – and I do own his collected works – he has not done research in the area of commutative algebra. There must be a story there...

i.e., the k-points on A are precisely the $\overline{k}$-points on $A_{\overline{k}}$ which are pointwise fixed under the G-action. Moreover, for any subextension l of $\overline{k}$, the invariants by $\mathrm{Aut}(\overline{k}/l)$ give the solutions to $V(J)$ defined over the field extension l.

Exercise: Why is this not true if k isn't perfect? Give a corresponding statement. (Hint: you will want to use the words "purely inseparable extension".)

However, this does not quite tell us the relationship between the maximal ideals of A and those of $\overline{A}$. In particular, if $\mathfrak{m}$ is a maximal ideal of A, then $\overline{m}$ need **not** be maximal in $\overline{A}$. If this seems surprising, we had better consider a familiar special case:

Example: Let $k = \mathbb{R}$, $A = \mathbb{R}[t]$, so $\overline{A} = \mathbb{C}[t]$. The maximal ideals of A are all principal, generated by monic irreducible polynomials, which, as a consequence of the fundamental theorem of algebra, are either linear: $t - r$ for $r \in \mathbb{R}$, or irreducible quadratic: $(t - r_1)(t - r_2)$, where $r_2 = \overline{r_1}$ are conjugate complex numbers. The first case is that of an $\mathbb{R}$-point, and those ideals remain maximal upon base change to $\mathbb{C}$. But in the second case the polynomial splits into a product of two Galois conjugate linear polynomials. In particular, in the ring extension $\mathbb{C}[t]$ of $\mathbb{R}[t]$, both $(t - r_1)$ and $(t - r_2)$ are maximal ideals lying over the maximal ideal $(t - r_1)(t - r_2)$ of $\mathbb{R}[t]$. Thus the maximal ideals of $\mathbb{R}[t]$ correspond not to $\mathbb{C}$-points of $\mathbb{C}[t]$, but to $G = \mathrm{Aut}(\mathbb{C}/\mathbb{R})$-orbits of $\mathbb{C}$-points of $\mathbb{C}[t]$. In other words, $\mathrm{MaxSpec}\,\mathbb{R}[t]$ is, as a set, the **quotient** of $\mathrm{MaxSpec}\,\mathbb{C}[t]$ by the natural action of G.

This is exactly what happens in general:

Theorem 8. *Let A be an affine algebra over a field k. Let $\iota : A \to \overline{A} = A \otimes_k \overline{k}$.*
a) For every maximal ideal $\mathfrak{m}$ of A, the set $\mathcal{M}(\mathfrak{m})$ of maximal ideals $\mathcal{M}$ of $\overline{A}$ lying over $\mathfrak{m}$ – i.e., such that $\mathcal{M} \cap A = \mathfrak{m}$ – is finite and nonempty.
b) Moreover, the natural action of $G = \mathrm{Aut}(\overline{k}/k)$ on $\mathcal{M}(\mathfrak{m})$ is transitive. Thus $\mathrm{MaxSpec}(A) = G\backslash \mathrm{MaxSpec}(\overline{A})$.
c) If k is perfect, the size of the G-orbit on $\mathfrak{m} \in \mathrm{MaxSpec}$ is equal to the degree of the field extension of k generated by the coordinates in $\overline{k}^n$ of any $\mathcal{M}$ lying over $\mathfrak{m}$.

In order to encourage you to think about this important correspondence, we leave the proof as an exercise. But we will give a hint: the case of $A = k[t_1, \ldots, t_n]$ itself is handled by the following theorem from commutative algebra:

Theorem 9. *(Matsumura, Theorem 9.3, pp. 66) Let R be an integrally closed domain, K its fraction field, L/K a normal algebraic field extension, and S the integral closure of R in L. Then, for any $\mathfrak{p} \in \mathrm{Spec}\, R$, $G = \mathrm{Aut}(L/K)$ acts transitively on the set of all prime ideals of S lying over $\mathfrak{p}$.*

Note that because the extension S/R is integral, if $\mathfrak{p}$ is maximal, so is every prime ideal of S lying over $\mathfrak{p}$, and this set is always nonempty. (In this level of generality, the set of such primes may well be infinite, as it will be e.g. for the case $R = \mathbb{Z}$, $L = \overline{\mathbb{Q}}$.)

Now reduce to the case of quotients by a radical ideal J and apply the Hilbert-Jacobson Nullstellensatz.

The point of all this is that we can now speak of points on affine varieties both in terms of elements of $\overline{k}^n$ and (in a closely related but not identical way) in terms of maximal ideals of the coordinate ring.

Exercise: Let A be an affine k-algebra. Define an $\mathrm{Aut}(\overline{k}/k)$-action on the set of geometric irreducible components of $\overline{A}$.

Example: We return to the example of $\mathbb{R}[x,y]/(x^2+y^2)$. This is an integral affine curve embedded in the affine plane, whose only $\mathbb{R}$-point is $(0,0)$. All the other maximal ideals have residue field $\mathbb{C}$ so their Galois orbits have size 2. Indeed, the Galois action also interchanges the two irreducible components, which are a pair of Galois conjugate lines.

Note that the unique $\mathbb{R}$-rational point on this curve is a singular point. This is a general fact of $\mathbb{R}$-analytic geometry: by a version of the implicit function theorem, locally near any nonsingular $\mathbb{R}$-point of a real algebraic variety of positive dimension d, the variety has the structure of a d-dimensional real manifold: in particular, isolated nonsingular points are not possible. There is also a more algebraic explanation for the above phenomenon:

Exercise: Let k be a perfect field, and A an affine k-algebra. Suppose that the $G = \mathrm{Aut}(\overline{k}/k)$-action on the geometric irreducible components is fixed-point free. Show that any k-rational point of A is a singular (i.e., nonregular) point. Does this continue to hold if k is not perfect?