A Brief Introduction to Model Theory: Hilbert's Nullstellensatz

Leonardo Finzi

In this write-up, we synthesize the discussions in Chapters 1-3 of [Mar02] to give a proof of a version of Hilbert's (Strong) Nullstellensatz as in [Mar02]. Most of this discussion is self-contained, but for more examples and explanations see [Mar02].

I want to thank Dr. Thomas Gilton for his help in understanding the beginnings of model theory, and his support and encouragement in applying this beautiful subject to algebra and algebraic geometry.

1. Basic Notions in Model Theory

A *language* $\mathcal{L}$ is a tuple $(\mathcal{F}, \mathcal{R}, \mathcal{C})$, where $\mathcal{F}$ is a set of function symbols and $n_f \in \mathbb{Z}_{>0}$ for each $f \in \mathcal{F}$, where $\mathcal{R}$ is a set of relation symbols and $n_R \in \mathbb{Z}_{>0}$ for each $R \in \mathcal{R}$, and where $\mathcal{C}$ is a set of constant symbols. Using a language, we can rigorously develop the notion of a formula.

Definition 1.1. Let $\mathcal{L} = (\mathcal{F}, \mathcal{R}, \mathcal{C})$ be a language. The set of $\mathcal{L}$-*terms* is the smallest set $\mathcal{T}$ such that $\mathcal{C} \subseteq \mathcal{T}$, each variable symbol $v_i \in \mathcal{T}$ for $i > 0$, and if $t_1, \dots, t_{n_f} \in \mathcal{T}$ and $f \in \mathcal{F}$, then $f(t_1, \dots, t_{n_f}) \in \mathcal{T}$. We say ϕ is an *atomic* $\mathcal{L}$-*formula* if ϕ is $t_1 = t_2$ for $t_1, t_2 \in \mathcal{T}$ or $R(t_1, \dots, t_{n_R})$ for $R \in \mathcal{R}$ and $t_1, \dots, t_{n_R} \in \mathcal{T}$. The set of $\mathcal{L}$-formulas is the smallest set $\mathcal{W}$ containing the atomic formulas such that, if $\phi, \psi \in \mathcal{W}$, then

$$\neg\phi, (\phi \wedge \psi), (\phi \vee \psi), \exists v_i \phi, \forall v_i \phi \in \mathcal{W}.$$

A $\mathcal{L}$-formula is a *sentence* if it contains no free variables.

For example, consider the language $\mathcal{L}_g = \{*, e\}$. If ϕ is $\exists v_i (e * v_i = e)$ and ψ is $\exists v_i (e * v_i = v_j)$, then ϕ is an $\mathcal{L}_g$-sentence while ψ is an $\mathcal{L}_g$-formula. This is because, in ψ, the variable in v_j is free, in that there is no quantifier bounding it.

Definition 1.2. Let $\mathcal{L}$ be a language. An $\mathcal{L}$-*structure* is a tuple

$$\mathcal{M} = (M, f^{\mathcal{M}}, R^{\mathcal{M}}, c^{\mathcal{M}} : f \in \mathcal{F}, R \in \mathcal{R}, c \in \mathcal{C}).$$

If $\mathcal{M}, \mathcal{N}$ are $\mathcal{L}$-structures, then an $\mathcal{L}$-*embedding* $\mathcal{M} \to \mathcal{N}$ is an injective map $\eta : M \to N$ such that

(i) for all $f \in \mathcal{F}$, the following diagram commutes:

$$\begin{array}{ccc} M^{n_f} & \xrightarrow{f^{\mathcal{M}}} & M \\ \downarrow{\scriptstyle \eta^n} & & \downarrow{\scriptstyle \eta} \\ N^{n_f} & \xrightarrow[f^{\mathcal{N}}]{} & N \end{array}$$

(ii) for $R \in \mathcal{R}$ and $(a_1, \ldots, a_{m_R}) \in M^{m_R}$, we have $(a_1, \ldots, a_{m_R}) \in R^{\mathcal{M}}$ if and only if $(\eta(a_1), \ldots, \eta(a_{m_R})) \in R^{\mathcal{N}}$; and

(iii) for all $c \in \mathcal{C}$, $\eta(c^{\mathcal{M}}) = c^{\mathcal{N}}$.

If $\mathcal{M}, \mathcal{N}$ are $\mathcal{L}$-structures such that $M \subseteq N$ and the inclusion $M \to N$ is an $\mathcal{L}$-embedding, then $\mathcal{M}$ is a *substructure* of $\mathcal{N}$.

Certain $\mathcal{L}$-substructures will be of importance to us.

Definition 1.3. Let $\mathcal{L}$ be a language and $\mathcal{M}, \mathcal{N}$ be $\mathcal{L}$-structures such that $\mathcal{M} \subseteq \mathcal{N}$. We say $\mathcal{M}$ is an *elementary substructure* of $\mathcal{N}$ and write $\mathcal{M} \prec \mathcal{N}$ if for all $\mathcal{L}$-formulas $\phi(v_1, \ldots, v_n)$ and $(a_1, \ldots, a_n) \in M^n$,

$$\mathcal{M} \models \phi(a_1, \ldots, a_n) \iff \mathcal{N} \models \phi(a_1, \ldots, a_n).$$

Consider the language $\mathcal{L}_g$ again. Given a group G, letting $*^{\mathcal{G}}$ be the binary operation on G and $e^{\mathcal{G}}$ be the identity element of G, we see that $\mathcal{G} = (G, *^{\mathcal{G}}, e^{\mathcal{G}})$ is an $\mathcal{L}_g$-structure. But an arbitrary $\mathcal{L}_g$-structure need not be a group, as $\mathcal{N} = (\mathbb{N}, +, 0)$ is an $\mathcal{L}_g$-structure by interpreting $*^{\mathcal{N}} = +$ and $e^{\mathcal{N}} = 0$. In fact, any set equipped with a binary relation and a distinguished element will be an $\mathcal{L}_g$-structure.

The motivation behind $\mathcal{L}$-structures is to use them to interpret collections of formulas and sentences in first-order logic. In this way, $\mathcal{L}$-structures give "meaning" to $\mathcal{L}$-formulas.

Definition 1.4. Let $\mathcal{L}$ be a language and $\mathcal{M}$ an $\mathcal{L}$-structure. Let ϕ be a formula in free variables $\overline{v} = (v_1, \ldots, v_n)$ and let $\overline{a} = (a_1, \ldots, a_n) \in M^n$. We define $\mathcal{M} \models \phi(\overline{a})$ inductively:

(i) If ϕ is $t_1 = t_2$, then $\mathcal{M} \models \phi(\overline{a})$ if $t_1^{\mathcal{M}}(\overline{a}) = t_2^{\mathcal{M}}(\overline{a})$.

(ii) If ϕ is $R(t_1, \ldots, t_n)$, then $\mathcal{M} \models \phi(\overline{a})$ if $(t_1^{\mathcal{M}}(\overline{a}), \ldots, t_n^{\mathcal{M}}(\overline{a})) \in R^{\mathcal{M}}$.

(iii) If ϕ is $\neg\psi$, then $\mathcal{M} \models \phi(\overline{a})$ if $\mathcal{M} \not\models \psi(\overline{a})$.

(iv) If ϕ is $\psi \wedge \theta$, then $\mathcal{M} \models \phi(\overline{a})$ if $\mathcal{M} \models \psi(\overline{a})$ and $\mathcal{M} \models \theta(\overline{a})$.

(v) If ϕ is $\psi \vee \theta$, then $\mathcal{M} \models \phi(\overline{a})$ if $\mathcal{M} \models \psi(\overline{a})$ or $\mathcal{M} \models \theta(\overline{a})$.

(vi) If ϕ is $\exists v_j \psi(\overline{v}, v_j)$, then $\mathcal{M} \models \phi(\overline{a})$ if there exists $b \in M$ such that $\mathcal{M} \models \psi(\overline{a}, b)$.

(vii) If ϕ is $\forall v_j \phi(\overline{v}, v_j)$, then $\mathcal{M} \models \phi(\overline{a})$ if $\mathcal{M} \models \psi(\overline{a}, b)$ for all $b \in M$.

If $\mathcal{M} \models \phi(\overline{a})$, we say $\mathcal{M}$ *satisfies* $\phi(\overline{a})$.

As an example, consider the $\mathcal{L}_g$-sentence ϕ defined as

$$\exists v\,(v \neq e \wedge v * v = e)\,.$$

First, consider the $\mathcal{L}_g$-structure $\mathbb{Z}/2\mathbb{Z} = (\mathbb{Z}/2\mathbb{Z}, +, 0)$, naturally interpreting $*^{\mathbb{Z}/2\mathbb{Z}} = +$ and $e^{\mathbb{Z}/2\mathbb{Z}} = 0$. If $v^{\mathbb{Z}/2\mathbb{Z}} = 1$, then $\mathbb{Z}/2\mathbb{Z} \not\models v = e$ and $\mathbb{Z}/2\mathbb{Z} \models v * v = e$, so $\mathbb{Z}/2\mathbb{Z} \models \phi$. But, if we take the $\mathcal{L}_g$-structure $\mathbb{Z} = (\mathbb{Z}, +, 0)$ with the natural interpretations of $*$ and e, we see that $\mathbb{Z} \not\models \phi$ since there is no non-trivial $z \in \mathbb{Z}$ that is torsion.

Note that, in some sense, this notion of satisfaction is what makes a formula "true." Of course some formulas, like $\neg(v_1 = v_1)$, will never be satisfied and others, like $(v_1 = v_1)$, will always be satisfied, but other formulas, like $(v_1 = v_2)$ or ϕ as above, will be true depending on the $\mathcal{L}$-structure chosen.

So far, we have seen that languages can admit many different types of structures. We now introduce a way to reduce the scope of our formal language to specific structures we wish to consider.

Definition 1.5. Let $\mathcal{L}$ be a language. A *$\mathcal{L}$-theory* T is a set of $\mathcal{L}$-sentences ϕ. We say an $\mathcal{L}$-structure $\mathcal{M}$ *models* T, and write $\mathcal{M} \models T$ if for all $\phi \in T$, $\mathcal{M} \models \phi$. The theory T is *satisfiable* if it has a model. We write $T \models \phi$ if for all models $\mathcal{M}$ of T, $\mathcal{M} \models \phi$.

For example, let T be the $\mathcal{L}_g$-theory of groups, i.e., the theory axiomatized by:

$$\forall x \forall y \forall z (x * (y * z) = (x * y) * z = x * y * z)$$

$$\forall x (e * x = x * e = x)$$

$$\forall x \exists y (x * y = y * x = e)$$

We see that because T is defined by the axioms of a group and that all groups satisfy the axioms of a group, that if G is a group, $G \models T$. Hence, we may think of theories as a formal and structured way of defining axioms of certain structures.

2. Quantifier Elimination

We will now introduce an important property that is extremely convenient for theories about algebraic structures to have.

Definition 2.1. We say an $\mathcal{L}$-theory T has *quantifier elimination* if for all $\mathcal{L}$-formulas $\phi(v)$, there exists a quantifier-free $\mathcal{L}$-formula $\psi(v)$ such that $T \models \phi(v) \leftrightarrow \psi(v)$.

We already implicitly have some understanding of quantifier elimination. We know by the Fundamental Theorem of Algebra that for any polynomial $p \in \mathbb{C}[x]$ of degree n, that there exists exactly n roots of p in $\mathbb{C}$, counting multiplicities. We can view this as a quantifier-elimination statement.

Let $\mathcal{L}_r = \{+, -, \cdot, 0, 1\}$ and let $\mathbb{C}$ be the $\mathcal{L}_r$-theory of the complex numbers. If $\phi(a_0, a_1, \ldots, a_n)$ is of the form $\exists x(a_n x^n + \cdots + a_1 x + a_0 = 0)$, then

$$\mathbb{C} \models \phi(a_0, a_1, \ldots, a_n) \leftrightarrow \left(a_0 = 0 \vee \bigvee_{i=1}^{n} a_i \neq 0 \right).$$

We will soon see that $\mathbb{C}$ can do this elimination for *every* $\mathcal{L}_r$-formula.

2.1. **Tests for Quantifier Elimination**

We now briefly develop our main example of a theory with quantifier-elimination, using the following tests. We first state a test for quantifier elimination that is somewhat involved.

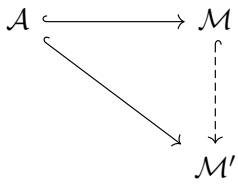
Theorem 2.2. Let T be an $\mathcal{L}$-theory. Suppose that for all quantifier-free formulas $\phi(\overline{v}, w)$, if $\mathcal{M}, \mathcal{N} \models T$, $\mathcal{A} \subseteq \mathcal{M} \cap \mathcal{N}$ is an $\mathcal{L}$-structure with $\overline{a} \in \mathcal{A}$, and $b \in \mathcal{M}$ such that $\mathcal{M} \models \phi(\overline{a}, b)$, then there is $c \in \mathcal{N}$ such that $\mathcal{N} \models \phi(\overline{a}, c)$. Then T has quantifier elimination.

Proof. For a proof, see [Mar02, pp. 73-75]. □

To understand the next, more practical test, we need some definitions. Note that for a theory T, the set $T_\forall$ consists of all universal consequences of T.

Definition 2.3. Let T be an $\mathcal{L}$-theory.

(i) The theory T has *algebraically prime models* if for any $\mathcal{A} \models T_\forall$, there is $\mathcal{M} \models T$ and an embedding $\mathcal{A} \hookrightarrow \mathcal{M}$ such that for all $\mathcal{M}' \models T$ and embeddings $\mathcal{A} \hookrightarrow \mathcal{M}'$, there exists an embedding $\mathcal{M} \hookrightarrow \mathcal{M}'$ such that the following diagram commutes:

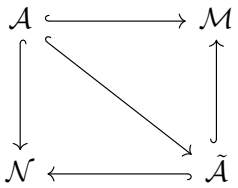


(ii) If $\mathcal{M}, \mathcal{N} \models T$ and $\mathcal{M} \subseteq \mathcal{N}$, we say $\mathcal{M}$ is *simply closed* in $\mathcal{N}$, and write $\mathcal{M} \prec_s \mathcal{N}$, if for any quantifier-free formula $\phi(\overline{v}, w)$ and $\overline{a} \in M^n$, if $\mathcal{N} \models \exists w\ \phi(\overline{a}, w)$, then $\mathcal{M} \models \exists w\ \phi(\overline{a}, w)$.

These two definitions allow us to state a more practical quantifier elimination test using Theorem 2.2.

Theorem 2.4. Suppose T is an $\mathcal{L}$-theory with algebraically prime models and such that for all $\mathcal{M}, \mathcal{N} \models T$ with $\mathcal{M} \subseteq \mathcal{N}$, $\mathcal{M} \prec_s \mathcal{N}$. Then T has quantifier elimination.

Proof. Let $\mathcal{M}, \mathcal{N} \models T$ and let $\mathcal{A} \subseteq \mathcal{M} \cap \mathcal{N}$. We may assume $\mathcal{A}$ is a model of $T_\forall$. Then there exists an algebraically prime model $\tilde{\mathcal{A}} \models T$ of $\mathcal{A}$ making the following diagram commute:



In particular, $\tilde{\mathcal{A}} \prec_s \mathcal{M}$ by assumption.

Let $\phi(\overline{v}, w)$ be a quantifier-free sentence such that $\mathcal{M} \models \phi(\overline{a}, b)$, for some $b \in M$. Then $\mathcal{M} \models \exists w\ \phi(\overline{a}, w)$. Since $\tilde{\mathcal{A}} \prec_s \mathcal{M}$, then $\tilde{\mathcal{A}} \models \exists w\ \phi(\overline{a}, w)$, and since $\tilde{\mathcal{A}} \hookrightarrow \mathcal{N}$, then $\mathcal{N} \models \exists w\ \phi(\overline{v}, w)$. In particular, there exists $c \in \mathcal{N}$ such that $\mathcal{N} \models \phi(\overline{a}, c)$. Thus, by Theorem 2.2, T has quantifier elimination. □

2.2. Algebraically Closed Fields

We turn our attention towards the theory of algebraically closed fields.

Definition 2.5. We define the $\mathcal{L}_r$-theory ACF to be the *theory of algebraically closed fields*, defined by the following axioms:

(i) $\forall x \forall y \forall z (x + (y + z) = (x + y) + z = x + y + z)$

(ii) $\forall x (0 + x = x + 0 = x)$

(iii) $\forall x \exists y (x + y = y + x = 0)$

(iv) $\forall x \forall y (x + y = y + x)$

(v) $\forall x \forall y \forall z (x - y = z \leftrightarrow x = y + z)$

(vi) $\forall x (x \cdot 0 = 0 \cdot x = 0)$

(vii) $\forall x \forall y \forall z (x \cdot (y \cdot z) = (x \cdot y) \cdot z)$

(viii) $\forall x (x \cdot 1 = 1 \cdot x = x)$

(ix) $\forall x \forall y \forall z (x \cdot (y + z) = (x \cdot y) + (x \cdot z))$

(x) $\forall x \forall y \forall z ((x + y) \cdot z = (x \cdot z) + (y \cdot z))$

(xi) $\forall x \forall y (x \cdot y = y \cdot x)$

(xii) $\forall x (x \neq 0 \rightarrow \exists y (x \cdot y = 1))$

(xiii) $\forall a_0 \ldots \forall a_{n_1} \exists x \left(x^n + \sum_{i=0}^{n-1} a_i x^i = 0 \right)$ for all $n > 0$

We note that (i)-(xi) defines the *theory of (commutative) rings* and (i)-(xii) defines the *theory of fields*.

We now show that ACF has quantifier elimination. We first want to understand $\text{ACF}_\forall$. Note that for a field F, we write $\overline{F}$ as the algebraic closure of F.

Lemma 2.6. $\text{ACF}_\forall$ is the theory of integral domains.

Proof. On the one hand, the theory of integral domains is defined by the universal axioms (i)-(xi) and

$$\forall x (x \neq 0 \rightarrow \forall y (y \neq 0 \rightarrow x \cdot y \neq 0)),$$

the last of which is a universal consequence of ACF. On the other hand, every integral domain D can be embedded inside $\overline{\text{Frac}(D)} \models \text{ACF}$ by the map $d \mapsto [(d, 1)]$, where $[-]$ denotes a fractional equivalence class. Hence, every integral domain is a subring of an algebraically closed field, and since universal sentences are preserved under substructure [Mar02, Proposition 1.1.8], if $\text{ACF}_\forall \models \phi$, then so does the theory of integral domains. □

The following lemma will be useful in our proof of quantifier-elimination in ACF.

Lemma 2.7. Let F be a field and $K \models \text{ACF}$. Suppose there is a non-trivial ring homomorphism $\phi : F \to K$. Then ϕ can be extended to a non-trivial ring homomorphism $\Phi : \overline{F} \to K$. In particular, Φ is an embedding.

Proof. Let F be a field and $K \models \text{ACF}$ such that there exists a non-trivial ring homomorphism $\phi : F \to K$. Let $F \subseteq F' = F(\alpha)$ be a finite algebraic field extension, for some $\alpha \in \overline{F}$. Define $\phi' : F' \to K$ such that $\phi'|_F = \phi$ and $\phi'(\alpha)$ is a root of the polynomial $\phi(f(x))$, where f is the minimal polynomial of α over F. Then ϕ' is a non-trivial ring homomorphism extending ϕ to F'.

Suppose $F = F_0 \subseteq F_1 \subseteq F_2 \subseteq \dots$ is a chain of finite field extensions of subfields of $\overline{F}$ and $\phi = \phi_0 \subseteq \phi_1 \subseteq \phi_2 \subseteq \dots$ is a chain of corresponding ring homomorphisms from $F_i \to K$. Consider the field

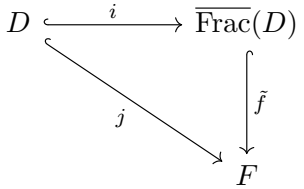
$$F^* = \bigcup_{i \in I} F_i$$

and the ring homomorphism $\phi^* : F^* \to K$ defined as $\phi^*(x) = \phi_i(x)$, where $x \in F_i$. Then (F^*, ϕ^*) is the upper bound of the chain $(F_0, \phi_0) \subseteq (F_1, \phi_1) \subseteq (F_2, \phi_2) \subseteq \dots$. By Zorn's Lemma, there exists a maximal element $(\tilde{F}, \Phi)$. Certainly $\tilde{F} \subseteq \overline{F}$.

Suppose there is $\alpha \in \overline{F} \setminus \tilde{F}$. Then $\tilde{F}(\alpha) \supset \tilde{F}$. So, as before, we can extend Φ to a strictly larger ring homomorphism $\Phi' : \tilde{F}(\alpha) \to K$, which is a contradiction to the maximality of Φ. Thus, $\tilde{F} = \overline{F}$. Since $\overline{F}$ is a field and Φ is non-trivial, then Φ is an injection. □

Theorem 2.8. ACF has quantifier elimination.

Proof. We will prove this using Theorem 2.4. If $D \models \text{ACF}_\forall$, then D is an integral domain by Lemma 2.6. Then $D \subseteq \text{Frac}(D) \subseteq \overline{\text{Frac}}(D)$, and D embeds into $\text{Frac}(D)$ (and hence $\overline{\text{Frac}}(D) \models \text{ACF}$) by the map $i : d \mapsto [(d, 1)]$, with the brackets denoting fractional equivalence class. Let $F \models \text{ACF}$ be such that D can be embedded in F through the map $j : D \to F$. Then the map $f : \text{Frac}(D) \to F$ taking $[(d, n)] \mapsto j(d)j(n)^{-1}$ is an injection because $\text{Frac}(D)$ is a field. Since f is an injection and $\text{Frac}(D) \subseteq \overline{\text{Frac}}(D)$, then by Lemma 2.7, f extends to an embedding $\tilde{f} : \overline{\text{Frac}}(D) \to F$, and the following diagram commutes:



Hence, ACF has algebraically prime models.

Let $F \subseteq K$ be models of ACF. We will show that $F \prec_s K$. Let $\phi(\overline{v}, w)$ be a quantifier-free formula and let $\overline{a} \in F^n$ such that $K \models \exists w \phi(\overline{a}, w)$. Then, for some $b \in K$, $K \models \phi(\overline{a}, k)$. We may assume that $\phi(\overline{v}, w)$ is the Boolean combination of atomic formulas. Since atomic formulas in $\mathcal{L}_r$ are of the form $p(\overline{v}, w) = 0$ for some $p \in \mathbb{Z}[x_1, \dots, x_\ell]$, then $\phi(\overline{a}, b)$ is of the form

$$\bigwedge_{i=1}^{n} p_i(\overline{a}, b) = 0 \wedge \bigwedge_{j=1}^{m} q_j(\overline{a}, b) \neq 0$$

for $p_i, q_j \in \mathbb{Z}[a_1, \ldots, a_n, b]$. In fact, since $\overline{a} \in F$, we may view $p_i, q_i \in F[b]$ and write $\phi(\overline{a}, b)$ as

$$\bigwedge_{i=1}^{n} p_i(b) = 0 \wedge \bigwedge_{j=1}^{m} q_j(b) \neq 0.$$

If any of the p_i are non-zero, then $b \in K$ would be algebraic over F, and $b \in F$ since F is algebraically closed. So, we may assume $\phi(\overline{a}, b)$ is of the form

$$\bigwedge_{j=1}^{m} q_j(b) \neq 0.$$

But as the degree of each q_j is finite, $q_j(w) = 0$ for finitely-many $w \in F$. Since $F \models \mathrm{ACF}$, it is infinite, so there are infinitely-many choices of $w \in F$ such that $q_j(w) \neq 0$. Hence, $F \models \exists w\ \phi(\overline{a}, w)$ and $F \prec_s K$. Thus, by Theorem 2.4, ACF has quantifier elimination. □

As a final remark on ACF, we will describe quantifier elimination in a slightly different away, which will make things easier for us later.

Definition 2.9. We say a theory T is *model-complete* if $\mathcal{M} \prec \mathcal{N}$ whenever $\mathcal{M}, \mathcal{N} \models T$ with $\mathcal{M} \subseteq \mathcal{N}$.

Proposition 2.10. If an $\mathcal{L}$-theory T has quantifier elimination, then it is model complete.

Proof. Let $\phi(v)$ be a $\mathcal{L}$-formula. Since T has quantifier-elimination, there exists an $\mathcal{L}$-formula $\psi(v)$ such that

$$T \models \phi(v) \leftrightarrow \psi(v).$$

Hence, if $\mathcal{M} \subseteq \mathcal{N}$ are models of T, then by [Mar02, Proposition 1.1.8],

$$\mathcal{M} \models \phi(v) \iff \mathcal{M} \models \psi(v) \iff \mathcal{N} \models \psi(v) \iff \mathcal{N} \models \phi(v).$$

Thus, $\mathcal{M} \prec \mathcal{N}$. □

Corollary 2.11. ACF is model-complete.

3. Hilbert's Nullstellensatz

We will now turn our attention towards Hilbert's Nullstellensatz. Although a brief definitional detour will be taken, we will closely utilize the fact that ACF has quantifier elimination to present a clear proof.

Definition 3.1. Let k be a field.

- For $S \subseteq k[x_1, \ldots, x_n]$, we define
$$V(S) = \{a \in k^n : \forall p \in S,\ p(a) = 0\}.$$
- For $Y \subseteq k^n$, we define
$$I(Y) = \{f \in k[x_1, \ldots, x_n] : \forall y \in Y,\ f(y) = 0\}.$$

Definition 3.2. Let $X \subseteq k^n$. We say X is *Zariski closed* in k^n if there exists $S \subseteq k[x_1, \ldots, x_n]$ such that $X = V(S)$.

We have some fascinating facts about these objects.

Proposition 3.3. Let k be a field.

(i) For $X \subseteq k^n$, $I(X)$ is a radical ideal, that is, if $f^n \in I(X)$ for some $n > 0$, then $f \in I(X)$.

(ii) If $X \subseteq k^n$ is Zariski closed, then $X = V(I(X))$.

(iii) If $X, Y \subseteq k^n$ are Zariski closed with $X \subseteq Y$, then $I(Y) \subseteq I(X)$.

(iv) If $X, Y \subseteq k^n$ are Zariski closed, then

$$X \cup Y = V(I(X) \cap I(Y)) \qquad X \cap Y = V(I(X) + I(Y)).$$

Proof. For a proof, see [Mar02, Lemma 3.2.4]. □

Consider the poset category $\mathtt{Top}_{\mathtt{Zar}}(k^n)$ of Zariski closed sets of k^n and the poset category $\mathtt{Rad}(k[x_1, \ldots, x_n])$ of radical ideals of $k[x_1, \ldots, x_n]$. Then Proposition 3.3 boils down to the statement that

$$\begin{aligned} \mathfrak{I} : \mathtt{Top}_{\mathtt{Zar}}(k^n) &\longrightarrow \mathtt{Rad}(k[x_1, \ldots, x_n]) \\ X &\longmapsto I(X) \\ (X \subseteq Y) &\longmapsto (I(Y) \subseteq I(X)) \end{aligned}$$

is a contravariant functor, i.e., an "arrow-inverting" functor. We will show that $\mathfrak{I}$ is an antiequivalence of categories, that is, that there is an inverse contravariant functor $\mathfrak{V} : \mathtt{Rad}(k[x_1, \ldots, x_n]) \to \mathtt{Top}_{\mathtt{Zar}}(k^n)$.

Theorem 3.4 (Primary Decomposition Theorem)**.** If $I \subset k[x_1, \ldots, x_n]$ is a radical ideal, then there are prime ideals $\mathfrak{p}_1, \ldots, \mathfrak{p}_m \supset I$ such that $I = \bigcap_{i=1}^m \mathfrak{p}_i$ and $I \neq \cup_{j \in J} \mathfrak{p}_j$, for any $J \subset \{1, \ldots, m\}$. Additionally, if $\mathfrak{q}_1, \ldots, \mathfrak{q}_n$ is another set of prime ideals with these properties, then $n = m$ and $\{\mathfrak{q}_1, \ldots, \mathfrak{q}_m\} = \{\mathfrak{p}_1, \ldots, \mathfrak{p}_m\}$.

Proof. For more information on this theorem, see Chapter 4 of [AM69]. □

Theorem 3.5. Let $k \models \mathrm{ACF}$. Suppose $I, J \subseteq k[x_1, \ldots, x_n]$ are radical ideals such that $I \subset J$. Then $V(J) \subset V(I)$.

Proof. Let $k \models \mathrm{ACF}$ and $p \in J \setminus I \neq \emptyset$. Then there exists a prime ideal $\mathfrak{p} \supset I$ such that $p \notin \mathfrak{p}$. Indeed, if no such prime ideal existed, then that means that for all $\mathfrak{p} \supset I$, $p \in \mathfrak{p}$. By Theorem 3.4, $p \in I$, which is a contradiction. Since $\mathfrak{p}$ is prime, then $k' = k[x_1, \ldots, x_n]/\mathfrak{p}$ is an integral domain, so $F = \overline{\mathrm{Frac}(k')} \models \mathrm{ACF}$.

Since $k[x_1, \ldots, x_n]$ is Noetherian, suppose $\mathfrak{p}$ is generated by $q_1, \ldots, q_m \in k[x_1, \ldots, x_n]$. Each $q_i \in \mathfrak{p}$ and $p \notin \mathfrak{p}$, so by interpreting a_i as $[x_i] \in F$, we have that

$$F \models \bigwedge_{i=1}^m q_i(\overline{a}) = 0 \wedge p(\overline{a}) \neq 0,$$

so

$$F \models \exists \overline{w} \left(\bigwedge_{i=1}^{m} q_i(\overline{w}) = 0 \wedge p(\overline{w}) \neq 0 \right).$$

By the model-completeness of ACF, since $k \subset F$ and $k \models \text{ACF}$,

$$k \models \exists \overline{w} \left(\bigwedge_{i=1}^{m} q_i(\overline{w}) = 0 \wedge p(\overline{w}) \neq 0 \right).$$

Hence, there is some $\overline{b} \in k$ such that $q_1(\overline{b}) = \cdots = q_m(\overline{b}) = 0$ and $p(\overline{b}) \neq 0$. In particular, $\overline{b} \in V(\mathfrak{p})$ but $\overline{b} \notin V(J)$. Since $I \subset \mathfrak{p}$, then $V(\mathfrak{p}) \subseteq V(I)$ and hence $V(I) \setminus V(J) \neq \emptyset$. Thus, $V(J) \subset V(I)$. □

Corollary 3.6 (Hilbert's Strong Nullstellensatz)**.** If $J \subseteq k[x_1, \ldots, x_n]$ is a radical ideal, then $J = I(V(J))$.

Proof. Suppose $p \in J$. Then for all $x \in V(J)$, $p(x) = 0$. Hence, $p \in I(V(J))$ and $J \subseteq I(V(J))$. On the other hand, since $V(J)$ is Zariski closed, then $V(J) = V(I(V(J)))$, and in particular, $V(J) \subseteq V(I(V(J)))$. By Theorem 3.5, we see that $I(V(J)) \subseteq J$. Thus, $J = I(V(J))$. □

Considering the functor

$$\begin{aligned} \mathfrak{V} : \mathtt{Rad}(k[x_1, \ldots, x_n]) &\longrightarrow \mathtt{Top}_{\mathtt{Zar}}(k^n)^{\mathtt{op}} \\ J &\longmapsto V(J) \\ (J \subseteq I) &\longmapsto (V(I) \subseteq V(J)) \end{aligned}$$

we see that the Nullstellensatz and Proposition 3.3(ii) tells us that the functors $\mathfrak{I}$ and $\mathfrak{V}$ are inverses, giving us that the categories $\mathtt{Top}_{\mathtt{Zar}}(k^n)$ and $\mathtt{Rad}(k[x_1, \ldots, x_n])$ are antiequivalent, or equivalently, that the Zariski closed sets in k^n are dual to the radical ideals of $k[x_1, \ldots, x_n]$.

References

[AM69] Michael F. Atiyah and I. G. Macdonald. Introduction to commutative algebra. Reading, Mass.-Menlo Park, Calif.-London-Don Mills, Ont.: Addison-Wesley Publishing Company (1969)., 1969.

[Mar02] David Marker. *Model theory: An introduction*, volume 217 of *Grad. Texts Math.* New York, NY: Springer, 2002.