CHAPTER 1

Tests

1. Tests for quantifier elimination

We have not been overly explicit about this, but for what follows it is important to have $\bot$ or $\top$ among one's atomic sentences.

DEFINITION 1.1. A theory T in a language L has *quantifier elimination* if for any L-formula $\varphi(x_1, \ldots, x_n)$ there is quantifier-free L-formula $\psi(x_1, \ldots, x_n)$ such that

$$T \models \varphi \leftrightarrow \psi.$$

We give two tests for quantifier elimination. The first one is simple:

DEFINITION 1.2. A *literal* is an atomic formula or a negated atomic formula. A formula will be called *primitive* if it is of the form

$$\exists x\, \varphi(\underline{y}, x)$$

where φ is a conjunction of literals.

PROPOSITION 1.3. *A theory T has quantifier elimination if and only if any primitive formula is equivalent over T to a quantifier-free formula.*

PROOF. Suppose every primitive formula is equivalent over T to a quantifier-free formula. Then every formula of the form

$$\exists x\, \varphi(\underline{y}, x)$$

with φ quantifier-free is also equivalent to a quantifier-free formula: for we can write $\varphi(\underline{y}, x)$ in disjunctive normal form, that is, as a disjunction $\bigvee_i \varphi_i(\underline{y}, x)$, where each $\varphi_i(\underline{y}, x)$ is a conjunctions of literals. Then we can push the disjunction through the existential quantifier, using that

$$\exists x \bigvee_i \varphi_i(\underline{y}, x) \leftrightarrow \bigvee_i \exists x\, \varphi_i(\underline{y}, x),$$

so that we are left with a disjunction of primitive formulas, which is equivalent to a quantifier-free formula, by assumption.

Now let φ be an arbitrary formula. We can rewrite φ into an equivalent formula using only $\neg$, $\wedge$ and $\exists$, and then, working inside out, eliminate all the existential quantifiers using the previous observation. $\square$

The second is a bit more complicated, but generally easier to apply. We need the following notion:

Definition 1.4. Let M and N be models. A *local isomorphism* is a map

$$f\colon \{m_1, \dots, m_n\} \subseteq M \to N$$

such that

$$M \models \varphi(m_1, \dots, m_n) \Leftrightarrow N \models \varphi(f(m_1), \dots, f(m_n))$$

holds for all quantifier-free formulas φ. (Note that this is equivalent to it holding for all atomic formulas.)

Theorem 1.5. *A theory T has quantifier elimination if and only if, given*

(1) *two models M and N of T, where N is ω-saturated,*
(2) *a local isomorphism $f\colon \{a_1, \dots, a_n\} \subseteq M \to N$, and*
(3) *an element $m \in M$,*

there is a local isomorphism $g\colon \{a_1, \dots, a_n, m\} \to N$ which extends f.

Proof. Necessity is clear: if T has quantifier elimination, then any local isomorphism is an elementary map, and we have

$$(M, a_1, \dots, a_n) \equiv (N, f(a_1), \dots, f(a_n)).$$

The type $p(x) = \mathrm{tp}_{(M,a_1,\dots,a_n)}(m)$ is satisfied in $(M, a_1, \dots, a_n)$, so it is finitely satisfiable in $(N, f(a_1), \dots, f(a_n))$. Because the latter is ω-saturated, we can find an element $n \in N$ realizing this type and we can extend f by putting $g(m) = n$.

Conversely, let L be the language of T and suppose $\exists x\, \varphi(\underline{y}, x)$ is a primitive formula not equivalent over T to a quantifier-free formula in L. Extend the language with constants $\underline{c}$ and work in the extended language. Now let T_0 be the collection of all quantifier-free sentences which are a consequence over T of $\neg\exists x\, \varphi(\underline{c}, x)$. Then the union of T, T_0 and $\exists y\, \varphi(\underline{c}, y)$ has a model M.

Next, consider T_1, which consists of the theory T, all quantifier-free sentences in the extended language which are true in M, as well as the sentence $\neg\exists y\, \varphi(\underline{c}, y)$. This theory T_1 is consistent: for if not, there would be a quantifier-free sentence $\psi(\underline{c})$ which is false in M and and which is a consequence of $\neg\exists x\, \varphi(\underline{c}, x)$ over T. But such a sentence must belong to T_0 and therefore be true in M. Contradiction!

So T_1 has a model N and we may assume that N is ω-saturated. Now let f be the map which sends the interpretation of c_i in M to its interpretation in N and let m be such that $M \models \varphi(\underline{c}, m)$. This f is a local isomorphism, but cannot be extended to one whose domain includes m, because $\exists y\, \varphi(\underline{c}, y)$ fails in N. □

2. Tests for completeness

Theorem 1.6. *(Vaught's Test) If a theory T in a language L is consistent, has no finite models and is λ-categorical for some $\lambda \geq |L|$, then T is complete.*

Proof. If T were not complete, there would a sentence ψ such that neither ψ nor $\neg\psi$ would follow from T. But then both $T \cup \{\psi\}$ and $T \cup \{\neg\psi\}$ would have infinite models. Since $\lambda \geq |L|$, both would actually have models of cardinality λ by the theorems of Skolem and Löwenheim. But these cannot be isomorphic, because they are not elementarily equivalent, contradicting the λ-categoricity of T. □

THEOREM 1.7. *If a theory T has quantifier elimination and there is a model M of T that can be embedded into every other model of T, then T is complete.*

PROOF. If N is any model of T, then M can be embedded into it. So M and N witness the same quantifier-free formulas with parameters from M. But since T has quantifier elimination, this implies that the same is true for *all* formulas with parameters from M. So the embedding is elementary and M and N are elementarily equivalent. Hence all models of T are elementary equivalent and so T must be complete. □

3. Tests for ω-saturation

THEOREM 1.8. *Let T be a nice theory. A model M of T is ω-saturated if and only if*

(1) *every countable model of T embeds elementarily into M, and*
(2) *M is ω-homogeneous.*

PROOF. We have already shown necessity, so we now show that (1) and (2) imply that M is ω-saturated. To that purpose let $X = \{m_1, \dots, m_k\}$ be a finite set of parameters from M and p be a 1-type over X which is finitely satisfiable in M. Then p can be realized in an elementary extension B of M, by an element b say, and by the downward Löwenheim-Skolem Theorem there is a countable elementary substructure A of B such that A contains both X and b.

Now, by assumption (1) there is an elementary embedding $f\colon A \to M$. Write $n_i = f(m_i)$, $a = f(b)$ and $Y = \{n_1, \dots, n_k\}$. Since f is an elementary embedding, we see that $g\colon Y \to X\colon n_i \mapsto m_i$ is an elementary map, so using that M is ω-homogeneous, we see that g can be extended to an elementary map h whose domain includes a. Then $m = h(a)$ realizes p, for we have

$$\begin{aligned} M \models \varphi(m_1, \dots, m_k, m) \quad &\Leftrightarrow \quad M \models \varphi(h(n_1), \dots, h(n_k), h(a)) \\ &\Leftrightarrow \quad M \models \varphi(n_1, \dots, n_k, a) \\ &\Leftrightarrow \quad M \models \varphi(f(m_1), \dots, f(m_k), f(b)) \\ &\Leftrightarrow \quad A \models \varphi(m_1, \dots, m_k, b) \\ &\Leftrightarrow \quad B \models \varphi(m_1, \dots, m_k, b) \\ &\Leftrightarrow \quad \varphi(m_1, \dots, m_k, x) \in p(x) \end{aligned}$$

for any formula $\varphi(\underline{y}, x)$. □

COROLLARY 1.9. *Suppose T is a nice theory with quantifier elimination and suppose M is a model of T such that:*

(1) *every countable model of T embeds into M, and*
(2) *if $a_1, \dots, a_n, m$ are elements from M and $f\colon \{a_1, \dots, a_n\} \to M$ is a local isomorphism, then f can be extended to a local isomorphism whose domain includes m.*

Then M is ω-saturated.

PROOF. If T has quantifier elimination then any embedding between models of T is elementary and any local isomorphism between models of T is an elementary map. So this follows from the previous theorem. □

Corollary 1.10. *Let κ be an infinite cardinal and suppose T is a κ-categorical theory in a countable language. If M is an ω-homogeneous model of cardinality κ, then M is ω-saturated.*

Proof. Since T must be a nice theory, it suffices to show that every countable model can be embedded into M. But if N_0 is a countable model of T then it embeds elementarily into a model of N_1 of cardinality κ, by the upward Löwenheim-Skolem theorem. But then $N_1 \cong M$, so that N_0 embeds elementarily into M as well. $\square$

CHAPTER 2

Examples

1. Dense linear orders

The theory DLO of dense linear orders without endpoints is the theory in the language $<$ saying that:

(1) $<$ defines an ordering: if $x < y$ then not $x = y$ and not $y < x$, and if $x < y$ and $y < z$ then $x < z$.
(2) The order $<$ is linear: $x < y$ or $x = y$ or $y < x$.
(3) It is dense: this says that $x < y$ implies that there is a z with $x < z < y$.
(4) It has no endpoints: for every x there are y and z such that $y < x < z$.

Examples are $\mathbb{Q}$ and $\mathbb{R}$.

THEOREM 2.1. *The theory DLO is ω-categorical.*

PROOF. Let M and N be two countable dense linear orders without endpoints. Fix enumerations $M = \{m_0, m_1, \ldots\}$ and $N = \{n_0, n_1, \ldots\}$. We will construct an increasing sequence of local isomorphisms f_k from some subset of M to N such that m_i belongs to the domain of f_{2i} and n_i belongs to the codomain of f_{2i+1}. Then $f = \bigcup_i f_i$ will be the desired isomorphism between M and N. We start with $f_0 = \emptyset$.

So suppose $k + 1 = 2i$ and we have constructed f_j for all $j \leq k$ and we want to construct f_{k+1}. If m_i already belongs to the domain of f_k, we do not need to do anything and we put $f_{k+1} = f_k$. If not, then we determine the relative position of m_i to all m belonging to the domain of $\mathrm{dom}(f_k)$. There are only a few possibilities: (1) m_i is small than all of these, (2) m_i is bigger than all of these, or (3) m_i is in between two elements $m < m'$ in the domain and then we may choose for m and m' its nearest neighbours so that no other element from the domain is in between m and m'. In case (1) we choose for $f_{k+1}(m_i)$ an element strictly smaller than all the elements in the image of f_k, in case (2) an element strictly bigger than all the elements in the image of f_k and in case (3) an element strictly between $f(m)$ and $f(m')$. This is possible since N is a dense linear order without endpoints.

If $k + 1 = 2i + 1$, we argue in the same way in order to find a suitable preimage for n_i. $\square$

We see from the proof: if M is a countable dense linear order, then any local isomorphism from a subset of M to itself can be extended to an automorphism of the entire structure M. And since every n-type is realized in M, we see that the n-types in variables $x_1, \ldots, x_n$ correspond to possible ways to order the x_i (while allowing for some of them to coincide). In particular, there are only finitely many of them and each of them is generated by a single quantifier-free formula. From this it follows:

THEOREM 2.2. *The theory DLO has quantifier elimination.*

PROOF. Let $[\varphi]$ be the open corresponding to a formula φ. It consists of finitely many n-types, each of which is generated by a quantifier-free formula. So let $\psi_1, \ldots, \psi_k$ be the quantifier-free formulas generating the n-types belonging to $[\varphi]$. Then $DLO \models \varphi \leftrightarrow \psi_1 \vee \ldots \vee \psi_k$. □

In fact, this would also have followed easily from Theorem 1.5.

EXERCISE 1. Show that DLO is not λ-categorical for any $\lambda > \omega$.

2. Random graph

DEFINITION 2.3. By a *graph* we will mean a pair (V, E) where V is a non-empty set and E is a binary relation on V which is both symmetric and irreflexive. We will refer to the elements of V as the *vertices* and the elements of E as the *edges*. If xEy holds for two $x, y \in V$, we say that x and y are *adjacent*.

A graph (V, E) will be called *random* if for any two finite sets of vertices X and Y which are disjoint there is a vertex $v \notin X \cup Y$ which adjacent to all of the vertices in X and to none of the vertices in Y. We will write RG for the theory of random graphs.

THEOREM 2.4. *If one is given*

(1) *two random graphs A and B,*
(2) *a local isomorphism $f\colon A_0 \subseteq A \to B$, where A_0 is finite, and*
(3) *an element $a \in A$,*

then one can extend the local isomorphism f to one whose domain includes a.

PROOF. Clearly, we only need to consider the case where $a \notin A_0$. Then consider the following two finite subsets of B:

$$\begin{aligned} X &= \{\, f(a_0)\colon a_0 \in A_0, a_0 \text{ adjacent to } a \,\}, \\ Y &= \{\, f(a_0)\colon a_0 \in A_0, a_0 \text{ not adjacent to } a \,\}. \end{aligned}$$

These subsets are disjoint and B is a random graph, so there is an element $b \in B$ which is adjacent to all elements in X and none of the elements in Y. This means that we can extend the local isomorphism f by putting $f(a) = b$. □

COROLLARY 2.5. (1) *The theory RG is ω-categorical.*
(2) *Any model of RG is ω-saturated.*
(3) *The theory RG has quantifier elimination.*

3. Atomless Boolean algebras

DEFINITION 2.6. A *(bounded) lattice* L is a partial order in which every finite subset $A \subseteq L$ has a least upper bound (a *supremum* or *join*, written $\bigvee A$) and a greatest lower bound (an *infimum* or *meet*, written $\bigwedge A$). More concretely this means that L has a smallest element 0, a largest element 1 and that for any two elements $p, q \in L$ there are elements $p \wedge q$ and $p \vee q$ such that:

$$\begin{aligned} x \leq p \wedge q \quad &\Leftrightarrow \quad x \leq p \text{ and } x \leq q, \\ p \vee q \leq x \quad &\Leftrightarrow \quad p \leq x \text{ and } p \leq x. \end{aligned}$$

Exercise 2. Show that in any lattice $\wedge$ and $\vee$ are associative, commutative and idempotent (that is, $x \wedge x = x$ and $x \vee x = x$ hold). In addition, show that the absorbative laws $x = x \wedge (x \vee y)$ and $x = x \vee (x \wedge y)$ hold, as well as $0 \wedge x = 0$ and $1 \vee y = y$.

Exercise 3. Conversely, show that if L is a set equipped with two binary operations $\wedge$ and $\vee$ and in which there are elements $0, 1 \in L$ such that all the properties from the previous exercise hold, then there is a unique ordering on L turning L into a lattice. (Hint: observe that in a lattice we have $x \leq y$ iff $x = x \wedge y$ iff $y = x \vee y$.)

Definition 2.7. A lattice L is called *distributive* if both distributive laws

$$\begin{aligned} x \wedge (y \vee z) &= (x \wedge y) \vee (x \wedge z), \\ x \vee (y \wedge z) &= (x \vee y) \wedge (x \vee z) \end{aligned}$$

are satisfied. A distributive lattice L is called a *Boolean algebra* if for any element $x \in L$ there is an element $\neg x \in L$ (its *complement*) for which

$$x \wedge \neg x = 0 \quad \text{and} \quad x \vee \neg x = 1$$

hold.

Example 2.8. For any set X the powerset $\mathcal{P}(X)$ is a Boolean algebra with order given by inclusion, meets and joins given by intersection and union, complements given by set-theoretic complement and smallest and largest elements $\emptyset$ and X.

Example 2.9. If X is a topological space, then the clopens in X also form a Boolean algebra with the same operations as in the previous example.

Exercise 4. Show that in any lattice one distributive law implies the other.

Exercise 5. Let L be a distributive lattice and suppose $x \in L$ is a complemented element, meaning that there is an element $y \in L$ such that $x \wedge y = 0$ and $x \vee y = 1$. Show that for any other element $p \in L$, we have

$$x \wedge p = 0 \Longrightarrow p \leq y \quad \text{and} \quad x \vee p = 1 \Longrightarrow y \leq p.$$

Deduce that complements are unique.

Exercise 6. Show that if B is a Boolean algebra, then B^{op}, which is B with the order reversed, is a Boolean algebra as well. In fact, B and B^{op} are isomorphic with the isomorphism given by negating (taking complements). Deduce the De Morgan laws: $\neg(p \wedge q) = \neg p \vee \neg q$ and $\neg(p \vee q) = \neg p \wedge \neg q$.

For what follows we need to understand finitely generated Boolean algebras. Recall that a Boolean algebra B is finitely generated if there are elements $b_1, \ldots, b_n \in B$ such that there is no proper Boolean subalgebra of B also containing the elements $b_1, \ldots, b_n$.

Theorem 2.10. *Finitely generated Boolean algebras are finite.*

Proof. Suppose B is generated by $b_1, b_2, \ldots, b_n$. Let C be the collection of elements in B that can be written as "conjunctions" of the form $c_1 \wedge c_2 \wedge \ldots \wedge c_n$ where c_i is either b_i or its complement, and let D the collection of elements in B that can be written as "disjunctions" of elements in C. The collections C and D are finite, because they contains at most 2^n and $2^{(2^n)}$ elements, respectively. But D is a Boolean subalgebra of B, because it contains 0 (no disjuncts), 1 (all disjuncts) and is closed under disjunction (clear), conjunction (by the distributive laws) and negation (by the De Morgan laws). So $B = D$ is finite; in fact, it contains at most $2^{(2^n)}$ many elements. $\square$

So we need to understand finite Boolean algebras. But these are always of the form $\mathcal{P}(X)$, where X is finite. To show this, we need some definitions.

DEFINITION 2.11. An element a in a Boolean algebra B is called an *atom* if $a > 0$ and there are no elements strictly in between a and 0. A Boolean algebra in which for any element $x > 0$ there is an atom a such that $x \geq a$ is called *atomic*. A Boolean algebra in which there are no atoms is called *atomless*.

PROPOSITION 2.12. *Finite Boolean algebras are atomic.*

PROOF. Let B is a finite Boolean algebra. Suppose $x_0 \in B$ is an element different from 0 and there are no atoms a with $x_0 \geq a$. This means that x_0 itself is no atom, so there is an element x_1 with $x_0 > x_1 > 0$. Of course, x_1 cannot be atom, by our assumption on x_0, so there must be an element x_2 such that $x_0 > x_1 > x_2 > 0$. Continuing in this way we create an infinitely descending sequence of elements in B, which contradicts its finiteness. □

PROPOSITION 2.13. *If B is an atomic Boolean algebra and $x < y$, then there is an atom $a \in B$ which lies below y, but not below x.*

PROOF. If $x < y$, then $y \wedge \neg x \neq 0$ (for if $y \wedge \neg x = 0$, then $\neg x \leq \neg y$ and $x \geq y$ by the exercises). So there is an atom a with $y \wedge \neg x \geq a$. So we have $y \geq a$ and $\neg x \geq a$; but the latter implies that $x \not\geq a$, for if also $x \geq a$, then $0 = x \wedge \neg x \geq a$. □

THEOREM 2.14. *All finite Boolean algebras B are of the form $\mathcal{P}(X)$ for a finite set X. In fact, X can be chosen to be the collection of atoms in B.*

PROOF. Let B be a finite Boolean algebra and let A be its collection of atoms. Then we define maps $f\colon B \to \mathcal{P}(A)$ by sending $b \in B$ to the set $f(b) = \{a \in A\colon a \leq b\}$ and $g\colon \mathcal{P}(A) \to B$ by sending a set $X \subseteq A$ to $g(X) = \bigvee X$. It will suffice to prove that f and g are order preserving and each other's inverses (since all operations in a Boolean algebra are uniquely determined in terms of its order, any order isomorphism between Boolean algebras must be an isomorphism of Boolean algebras). That they are order preserving is clear, so we only check that they are each other's inverses.

So if b is an element in B and $X = \{a \in A\colon a \leq b\}$, then b is an upper bound for X, so $b \geq \bigvee X$. Here we must have equality: for if $b > \bigvee X$, then the previous two results imply that there is an atom a' such that $b \geq a'$ but not $\bigvee X \geq a'$. But the former implies that $a' \in X$ so we should have $\bigvee X \geq a'$ after all. Contradiction! We deduce $g(f(b)) = b$.

Conversely, let X be a set of atoms in B and $b = \bigvee X$. Clearly, all atoms in X are below b, but the converse is true as well: for suppose a' is an atom and $b \geq a'$. Then

$$0 < a' = (a' \wedge b) = a' \wedge \bigvee_{a \in X} a = \bigvee_{a \in X} (a' \wedge a).$$

So there must be an element $a \in X$ such that $a' \wedge a$ is not zero. But since a and a' are atoms and $a' \wedge a$ is below each of them, we must have $a = a \wedge a' = a'$. We deduce $f(g(X)) = X$, which finishes the proof. □

THEOREM 2.15. *The theory ABA of atomless Boolean algebras is ω-categorical.*

PROOF. Observe that atomless Boolean algebras have to be infinite (by Proposition 2.12) and that there is a countable and atomless Boolean algebra: look at the clopens in Cantor space.

Let A and B be two countable atomless Boolean algebras and fix enumerations $a_1, a_2, \dots$ of A and $b_1, b_2, \dots$ of B. Again, we will construct a sequence $(f_n)_{n \in \mathbb{N}}$ of local isomorphisms from A to B with a_i in the domain of f_{2i} and b_i in the codomain of f_{2i-1}. Put $f_0 = \emptyset$.

Now suppose f_k has been constructed for all $k < n$ and we want to build f_n. Write C for the Boolean subalgebra of A generated by $a_0, \dots, a_{n-1}$ and D for the Boolean subalgebra of B generated by $b_0, \dots, b_{n-1}$. The local isomorphism f_{n-1} induces an isomorphism $\overline{f}$ of Boolean algebras from C to D and without loss of generality we may assume that $a_0, \dots, a_{n-1}$ are the atoms of C and $b_0, \dots, b_{n-1}$ are the atoms of D and $f(a_i) = b_i$.

For any $x \in A$, there are three possibilities for $x \wedge a_i$: it can be 0, or a_i or something in between. Let us call the function which says for every i which of these three possibilities happens, the *profile* of x. Similarly, we can define the profile of elements $y \in B$, but then with respect to the b_i instead of the a_i.

The proof will be finished once I show:

(1) For any $x \in A$ there is a $y \in B$ which has the same profile, and vice versa.
(2) If $x \in A$ and $y \in B$ have the same profile, then the local isomorphism can be extended to one which sends x to y.

I will only sketch the argument: as for (1), let $I = \{i < n : x \wedge a_i = a_i\}$ and $J = \{j < n : 0 < (x \wedge a_j) < a_j\}$. For any $j \in J$ we consider b_j: since it is not an atom in B, we can choose an element $y_j \in B$ with $0 < y_j < b_j$.

Now put $y := \bigvee_{i \in I} b_i \wedge \bigvee_{j \in J} y_j$. Using that the b_i are atoms in D and we therefore have that $b_i \wedge b_j = 0$ whenever $i \neq j$, we see that y has the same profile as x.

As for (2): the crucial observation here is that if $J = \{j < n : 0 < (x \wedge a_j) < a_j\}$, then the atoms of the Boolean subalgebra generated by $a_0, \dots, a_{n-1}$ and x are the a_i with $i \in J^c$ together with $a_j \wedge x$ and $a_j \wedge \neg x$ for every $j \in J$. Sending these to b_i, $b_j \wedge y$ and $b_j \wedge \neg y$, respectively, we have a maps from atoms to atoms which extends uniquely to a map of Boolean algebras: this one extends the original map and sends x to y. □

THEOREM 2.16. *The theory of atomless Boolean algebras has quantifier elimination.*

PROOF. An n-type in variables $x_1, \dots, x_n$ should say what the profile of x_i is in terms of the atoms of the Boolean subalgebra generated by $x_1, \dots, x_{i-1}$: call this a sequence of profiles. I claim that a sequence of profiles completely determines the n-type: by this I mean that if $a_1, \dots, a_n$ is a tuple in a model A and $b_1, \dots, b_n$ is a tuple in a model B and they determine the same sequence of profiles, then they realize the same type. For by the downward Lowenheim-Skolem Theorem, we may assume that both A and B are countable, in which case the proof of the previous theorem implies that there is an isomorphism from A to B sending a_i to b_i. Since a sequence of profiles can be formulated using a single quantifier-free sentence, and there are only finitely many n-types, the theory ABA has quantifier elimination. □

Again, we could also have used Theorem 1.5.

EXERCISE 7. Show that all Boolean algebras of the form $\mathcal{P}(X)$ are atomic, but that there are atomic Boolean algebras which are not of this form.

EXERCISE 8. Not so easy: show that ABA is not λ-categorical for any $\lambda > \omega$.

4. Vector spaces

For a fixed field K, the language of K-vector spaces contains symbols $+$ and 0, for vector addition and the null vector, as well as unary operations f_k, one for every $k \in K$, for scalar multiplication with k. The theory IVS_K of infinite vector spaces over K expresses that $(+, 0)$ is an infinite Abelian group on which K acts as a set of scalars.

THEOREM 2.17. *The theory IVS_K is λ-categorical for all $\lambda > |K|$.*

PROOF. Because vector spaces are completely determined by their dimension and if V is a vector space over a field K of cardinality $\lambda > |K|$, then its dimension is λ. □

THEOREM 2.18. *The theory IVS_K has quantifier elimination.*

PROOF. We will use Theorem 1.5. So let V and W be two infinite K-vector spaces, where W is ω-saturated, and let $f\colon \{v_1, \ldots, v_n\} \to W$ be a local isomorphism. If $v \in V$, then there are two possibilities: v is a linear combination $k_1 v_1 + \ldots + k_n v_n$, in which case v should be sent to $k_1 f(v_1) + \ldots + k_n f(v_n)$. Or v is linearly independent from $v_1, \ldots, v_n$: in case K is finite, we use that W is infinite, and in case K is infinite, we use ω-saturation of W to find a vector $w \in W$ which is linearly independent from $f(v_1), \ldots, f(v_n)$. Then extend f by putting $f(v) = w$. □

5. Algebraically closed fields

Recall that a field K is called algebraically closed if every non-constant polynomial has a root in K. For convenience, we will only consider fields of characteristic 0 and only consider ACF_0, the theory of algebraically closed fields of characteristic 0.

5.1. Recap on fields. Consider an inclusion $K \subseteq L$ of fields. Recall that L can be considered as a K-vector space and that we write $[K\colon L]$ for its dimension.

PROPOSITION 2.19. *If we have two field extensions $K \subseteq L \subseteq M$, then $[M\colon K] = [M\colon L][L\colon K]$.*

If $K \subseteq L$ and $\xi \in L$, then there are two possibilities:

1. ξ is algebraic over K. This means that there is a polynomial $p(x)$ with coefficients from K such that $p(\xi) = 0$. In this case we can consider the monic polynomial $m(x) \in K[x]$ with $m(\xi) = 0$ which has least possible degree: this is called the *minimal polynomial* of ξ. This polynomial has to be irreducible and $K(\xi)$, the smallest subfield of L which contains both K and ξ, is isomorphic to $K[x]/(m(x))$. In this case $[K(\xi)\colon K]$ is finite.
2. ξ is transcendental over K. In this case $K(\xi)$ is isomorphic to the quotient field $K(x)$ and $[K(\xi)\colon K]$ is infinite.

An extension $K \subseteq L$ is called *algebraic* if all elements in L are algebraic over K. From Proposition 2.19 it follows that:

1. $K(\xi)$ is algebraic over K precisely when ξ is algebraic over K.
2. If $K \subseteq L$ and $L \subseteq M$ are two field extensions and they are both algebraic, then so is $K \subseteq M$.

5.2. Algebraic closure.

DEFINITION 2.20. If $K \subseteq L$ is a field extension, then L is an *algebraic closure* of K, if L is algebraic over K, but no proper extension of L is algebraic over K.

THEOREM 2.21. *Algebraic closures are algebraically closed.*

PROOF. Let L be the algebraic closure of K and $p(x)$ be a non-constant polynomial with coefficients from L without any roots in L. Without loss of generality we may assume that $p(x)$ is irreducible (otherwise replace $p(x)$ with one of its irreducible factors); but then $L[x]/(p(x))$ is a proper algebraic extension of L and K, which is a contradiction. □

THEOREM 2.22. *Every field K has an algebraic closure.*

PROOF. Let X the collection of algebraic field extensions of K and order by embedding of fields. We restrict attention to those fields which have the same cardinality as K and therefore X is a set (essentially). Clearly, every chain of embeddings has an upper bound in X, so X has a maximal element L. This field is an algebraic closure of X: for if $L \subset M$ is a proper extension of fields and $\xi \in M - L$, then ξ cannot be algebraic over K. For otherwise, $L \subset L(\xi) \in X$, contradicting maximality of L. □

THEOREM 2.23. *Algebraic closures are unique up to (non-unique) isomorphism.*

PROOF. By a back and forth argument. Let L and M be algebraic closures of K. Since L and M have the same (infinite) cardinality as K, which is κ say, we can fix enumerations $\{l_i : i \in \kappa\}$ and $\{m_i : i \in \kappa\}$ of L and M, respectively. By induction on $i \in \kappa$ we will construct an increasing sequence of isomorphisms $f_i \colon L_i \to M_i$ between subfields of L and M such that $\bigcup L_i = L$ and $\bigcup M_i = M$. We start by declaring f_0 to be isomorphism between the isomorphic copies of K inside L and M; and at limit stages we simply take the union.

If $i + 1 = 2j$, then look at the minimal polynomial $m(x) = a_n x^n + a_{n-1} x^{n-1} + \ldots + a_0$ of l_j over L_i: such a thing exists because L is algebraic over K and hence over L_i. Because M is algebraically closed, there exists a root $m \in M$ of the polynomial $n(x) = f_i(a_n)x^n + f_i(a_{n-1})x^{n-1} + \ldots + f(a_0)$; since f_i is an isomorphism, the polynomial $n(x)$ is irreducible over M_i and $n(x)$ must be the minimal polynomial of m over M_i. So we can extend the isomorphism by sending l_j to m:

$$f_{i+1} \colon L_i(l_j) \cong L_i[x]/(m(x)) \cong M_i[x]/(n(x)) \cong M_i(m).$$

If $i + 1 = 2j + 1$, then we can use a similar argument to show that the isomorphism f_i can be extended to one whose codomain includes m_j. □

5.3. Categoricity.

A similar argument shows:

THEOREM 2.24. *The theory ACF_0 is λ-categorical for any uncountable λ.*

PROOF. Let L and M be two algebraically closed fields of the same uncountable cardinality λ and fix enumerations $\{l_i : i \in \lambda\}$ and $\{m_i : i \in \lambda\}$ of L and M, respectively. By induction on $i \in \lambda$ we will construct an increasing sequence of isomorphisms $f_i \colon L_i \to M_i$ between subfields of L and M *of cardinality strictly less than* λ such that $\bigcup L_i = L$ and $\bigcup M_i = M$. We start by declaring f_0 to be isomorphism between the isomorphic copies of the rationals inside L and M; and at limit stages we simply take the union.

If $i+1=2j$, then there are two possibilities for l_j *vis-à-vis* L_i: it can either be algebraic or transcendental. If it is algebraic, we proceed as in the proof of the previous theorem. We look at the minimal polynomial $m(x) = a_n x^n + a_{n-1}x^{n-1} + \ldots + a_0$ of l_j over L_i and use that M is algebraically closed to find an element $m \in M$ with minimal polynomial $n(x) = f_i(a_n)x^n + f_i(a_{n-1})x^{n-1} + \ldots + f(a_0)$ over M_i. And we extend the isomorphism by sending l_j to m:

$$f_{i+1}\colon L_i(l_i) \cong L_i[x]/(m(x)) \cong M_i[x]/(n(x)) \cong M_i(m).$$

If, one the other hand, l_j is transcendental over L_i, we use the fact that $|M_i| < |M|$ to deduce that M also contains an element $m \in M$ which transcendental over M_i. And the isomorphism can be extended by sending l_j to m:

$$f_{i+1}\colon L_i(l_j) \cong L_i(x) \cong M_i(x) \cong M_i(m).$$

If $i+1 = 2j+1$, then we can use a similar argument to show that the isomorphism f_i can be extended to one whose codomain includes m_j. □

Note, however, that the theory ACF_0 is not ω-categorical: consider, for example, the algebraic closures of $\mathbb{Q}$ and $\mathbb{Q}(\pi)$.

Corollary 2.25. *The theory ACF_0 is complete.*

5.4. Quantifier elimination.

Theorem 2.26. *The theory ACF_0 has quantifier elimination.*

Proof. We use Theorem 1.5. So let L and M be two algebraically closed fields, where M in addition is ω-saturated. We assume we are given a local isomorphism $f\colon \{l_1, \ldots, l_n\} \to M$ and an element $l \in L$ and we want to extend the local isomorphism f to one whose domain includes l.

Because it is a local isomorphism, the map f extends to an embedding of fields

$$\overline{f}\colon \mathbb{Q}(l_1, \ldots, l_n) \to M.$$

If l is algebraic over $\mathbb{Q}(l_1, \ldots, l_n)$ with minimal polynomial

$$m(x) = a_k x^k + a_{k-1}x^{k-1} + \ldots + a_0,$$

then we send l to an element $m \in M$ whose minimal polynomial is

$$\overline{f}(a_k)x^k + \overline{f}(a_{k-1})x^{k-1} + \ldots + \overline{f}(a_0)$$

over $\mathrm{Im}(\overline{f})$. If l is transcendental over $L(l_1, \ldots, l_n)$, we use that M is ω-saturated to find an element $m \in M$ which is transcendental over $\overline{f}(l_1), \ldots, \overline{f}(l_n)$ and we send l to m. □

6. Real closed ordered fields

6.1. Ordered fields.

Definition 2.27. An *ordered field* is a field equipped with a linear order $\leq$ satisfying

(1) if $x \leq y$, then $x + z \leq y + z$,
(2) if $x \leq y$ and $0 \leq z$, then $xz \leq yz$.

Let us call elements x for which $x \geq 0$ *positive*; otherwise x is called negative. Note that if x is negative, then $x < 0$ and

$$-x = 0 - x \geq x - x = 0,$$

so $-x$ is positive. Using property (2) and the observation that $x^2 = (-x)^2$, it follows that $1 = 1^2$ is positive and also $2, 3, 4, \ldots$ are positive. But -1 is negative and hence ordered fields always have characteristic 0.

DEFINITION 2.28. If K is a field, then we call a subset $P \subseteq F$ a *positive cone*, if:

(1) P is closed under sums and products.
(2) $-1 \notin P$.
(3) for any x, either x or $-x$ belongs to P.

PROPOSITION 2.29. *If K is an ordered field, then the elements $x \in K$ satisfying $x \geq 0$ form a positive cone. Conversely, if P is a positive cone on a field K, then K can be ordered by putting $x \leq y$ iff $y - x \in P$.*

In ordered fields sums of squares have to be positive. In fact, we have:

PROPOSITION 2.30. *Let K be a field and $r \in K$. If both -1 and r cannot be written as a sum of squares, then K can be ordered in such a way that r becomes negative.*

PROOF. Let S be the collection of those elements in K that can be written as sums of squares. This set has the following properties:

(1) it is closed under sums and products,
(2) it contains all squares,
(3) and it does not contain -1.

Such a set is called a *semipositive cone*. We use two properties of such sets: first, if X is a semipositive cone and $s \in X - \{0\}$, then $(\frac{1}{s})^2 \in X$ and hence also $\frac{1}{s} \in X$. And if X is a semipositive cone and $s \notin X$, then $X - sX$ is also semipositive cone. For if there would be x_0, x_1 such that $x_0 - sx_1 = -1$, then $x_1 \neq 0$ and

$$s = \frac{1 + x_0}{x_1} \in X.$$

So put $Y := S - rS$. This is a semipositive cone, and, using Zorn's Lemma, we can extend Y to a maximal semipositive cone Y_{max}. Then Y_{max} is a positive cone, for if $x \notin Y_{max}$, then $-x \in Y_{max} - xY_{max} = Y_{max}$. □

6.2. Some analysis in ordered fields. Now suppose that K is an ordered field.

PROPOSITION 2.31. *Let $p(x) = x^d + a_{d-1}x^{d-1} + \ldots + a_0$ and $m = \max(|a_{d-1}|, \ldots, |a_0|) + 1$. Then all roots of $p(x)$ lie between $-m$ and m.*

PROOF. If $|x| \geq m$, then

$$|P(x) - x^d| \leq (|m| - 1)\,(|x|^{d-1} + |x|^{d-2} + \ldots + 1) \leq (|m| - 1)\,\frac{|x|^d - 1}{|x| - 1} \leq |x|^d - 1$$

so $P(x) \neq 0$. □

PROPOSITION 2.32. *If $p(x) \in K[x]$ and $p(0) > 0$, then there is an $\epsilon > 0$ such that $P(x) > 0$ for all $x \in [-\epsilon, +\epsilon]$.*

PROOF. Let $p(x) = a_d x^d + a_{d-1}x^{d-1} + \ldots + a_0$. Then put $m = \max(|a_d|, |a_{d-1}|, \ldots, |a_0|)$ and $\epsilon = \min(1, \frac{P(0)}{2md})$. Then $x \in [-\epsilon, +\epsilon]$ implies

$$\begin{aligned}
|p(x) - p(0)| &\leq |a_d x^d + a_{d-1}x^{d-1} + \ldots + a_0 - a_0| \\
&\leq m\epsilon^d + m\epsilon^{d-1} + \ldots + m\epsilon \\
&\leq md\epsilon \\
&\leq \frac{1}{2}p(0)
\end{aligned}$$

and hence $p(x) > 0$. □

PROPOSITION 2.33. *If $p'(a) > 0$, then there is an $\epsilon > 0$ such that $p(x) > p(a)$ for every $x \in (a, a+\epsilon]$ and $p(x) < p(a)$ for every $x \in [a-\epsilon, a)$.*

PROOF. Write $p(x) = (x-a)q(x)+p(a)$. Then $p'(x) = q(x)+(x-a)q'(x)$, so $q(a) = p'(a) > 0$. Then choose ϵ such that $q(x) > 0$ for all $x \in [a-\epsilon, a+\epsilon]$ using the previous result. □

6.3. Real closed ordered fields.

DEFINITION 2.34. An ordered field will be called *real closed* if it satisfies the intermediate value theorem for polynomials: if for any polynomial $P(x)$ and elements $a < b$ such that $P(a) < 0$ and $P(b) > 0$ there is an element $c \in (a, b)$ such that $P(c) = 0$.

For example, the field $\mathbb{R}$ is real closed, but $\mathbb{Q}$ is not.

PROPOSITION 2.35. *In a real closed field an element is positive iff it can be written as a square.*

PROOF. We already know that squares are positive. So suppose $a > 0$ and consider $p(x) = x^2 - a$. Then $p(a+1) = (a+1)^2 - a = a^2 + 2a + 1 - a = a^2 + a + 1 > 0$ and $p(0) < 0$, so there is an element r such that $p(r) = 0$ and hence $r^2 = a$. □

EXERCISE 9. Use this to say in the language of fields (without order!) that the field can be ordered in such a way that it becomes real closed.

THEOREM 2.36. *Let K be a real closed field and $p(x)$ be a polynomial over K. If $a < b \in K$ and $p'(x) > 0$ for all $x \in (a, b)$, then $p(a) < p(b)$.*

PROOF. First suppose that $p'(a) > 0$ and $p'(b) > 0$. Then we can use Proposition 2.33 to find c, d with $a < c < d < b$ such that $p(a) < p(c)$ and $p(d) < p(b)$. So if $p(a) \geq p(b)$, then $p(c) > p(b) > p(d)$ and there is an $e_0 \in (c, d)$ such that $p(e_0) = p(b)$. By repeating this argument for e_i and b instead of a and b we find for every $i \in \mathbb{N}$ an $e_{i+1} \in (e_i, b)$ such that $p(e_{i+1}) = p(b)$, contradicting the fact that a polynomial can have only finitely many zeros.

In the general case choose arbitrary c, d such that $a < c < d < b$. We have $p(c) < p(d)$ by the previous argument. In addition, we have $p(a) \leq p(c)$, for if $p(a) > p(c)$, then there is an $e \in (a, c)$ such that $p(e) > p(c)$ by continuity of p. But that again contradicts the previous argument. Similary, $p(c) \leq p(d)$, so $p(a) < p(b)$. □

COROLLARY 2.37. (Rolle's Theorem for real closed ordered fields) *Let K be a real closed ordered field and $p(x)$ be a polynomial over K. If $p(a) = p(b)$ for $a < b$, then there exists $c \in (a, b)$ with $P'(c) = 0$.*

Proof. For if $P'(c) \neq 0$ for all $c \in (a, b)$, then P' is either strictly positive or strictly negative on (a, b), by real closure. $\square$

6.4. Real closure.

Definition 2.38. Let $K \subseteq L$ be an order preserving embedding between ordered fields. L is a *real closure* of K, if L is algebraic over K and no ordered field properly extending L is algebraic over K.

Note, by the way, that an inclusion of ordered fields $K \subseteq L$ is order preserving iff it is order reflecting, because ordered fields are linearly ordered.

Theorem 2.39. *If L is a real closure of K, then L is real closed.*

Proof. Suppose there are polynomials in $L[x]$ for which the intermediate value theorem for polynomials fails. Let p be a counterexample of minimal degree: so the intermediate value theorem holds for polynomials in $L[x]$ with degree smaller than p, but there are $a < b \in L$ with $p(a) < 0$ and $p(b) > 0$ for which no $\xi \in (a, b)$ with $p(\xi) = 0$ exists.

In that case p has to be irreducible so $L[x]/(p(x))$ is a field extending L, still algebraic over K. So once we show that $L[x]/(p(x))$ can be ordered in a way which extends to the order on L, we have obtained our desired contradiction.

Let $A = \{x \in [a, b] : (\exists y \geq x)\, p(y) < 0\}$ and $B = [a, b] - B = \{x \in [a, b] : (\forall y \geq x)\, p(y) > 0\}$. Since polynomials are continuous, both A and B are open and have no greatest or least element, respectively. So if $q(x)$ is any non-zero polynomial, then q has only finitely many roots, so there are $a_0 \in A$ and $b_0 \in B$ such that q has no roots in the interval $[a_0, b_0]$. If $q(x)$ has a degree strictly smaller than $p(x)$, then the intermediate value theorem holds for $q(x)$ and $q(x)$ is either strictly positive or strictly negative on $[a_0, b_0]$. If the former holds we declare $q(x)$ positive. It is easy to see that this defines a positive cone on $L[x]/(p(x))$ extending the one on L. So we have our desired contradiction. $\square$

Theorem 2.40. *Real closures exist and are unique up to unique isomorphism.*

Proof. The existence of real closures follows from Zorn's Lemma: consider all ordered extensions of a field K which are still algebraic over K and all field embeddings between them which preserve the ordering. Since fields algebraic over K have the same infinite cardinality as K, this is essentially a set. Since chains have upper bounds given by unions, a maximal element must exist, which is a real closure of K.

Now suppose both L_0 and L_1 are real closures of an ordered field K. By Zorn's Lemma, again, there are subfields $K_0 \subseteq L_0$ and $K_1 \subseteq L_1$ between which there exists an order preserving isomorphism f which leaves K invariant and which is maximal with these properties. If either $L_0 - K_0$ or $L_1 - K_1$ is non-empty, then we may assume, without loss of generality, that there is an element $\xi \in L_0 - K_0$ with minimal polynomial $p(x)$ over K such that all other elements $\xi' \in L_i - K_i$ have a minimal polynomial over K whose degree is at least that of p.

Since p is minimal, we have $p'(\xi) \neq 0$, so p changes sign in ξ. Moreover, in L_1 and L_2 it holds that in between any two roots of $p(x)$ lies a root of $p'(x)$, by Rolle's Theorem. Since roots of $p'(x)$ have a minimal polynomial whose degree is strictly smaller than that of $p(x)$, these roots of $p'(x)$ lie already in K_0 and K_1. So for ξ there are three possibilities:

(1) ξ lies in between two roots of $p'(x)$, call them x_0 and x_1, and it is the only root lying in this interval. In that case p has different signs in x_0 and x_1. So the same applies to $f(x_0)$ and $f(x_1)$ and the polynomial p can have only one root in K_1 in between these points. Then ξ should be sent to this root.
(2) ξ is bigger than the largest root of $p'(x)$. Let x_0 be this largest root and let x_1 be a number in K bounding the zeros of p from above (using Proposition 2.31). Then again p changes sign between x_0 and x_1 and ξ should be sent to the unique root of p in K_1 between $f(x_0)$ and $f(x_1)$.
(3) ξ is smaller than the smallest root of $p'(x)$. Then the same argument as in (2) applies.

This determines a field isomorphism between $K(\xi) \cong K[x]/(p(x)) \cong K(\xi')$. The question now is why this field isomorphism should be order preserving. But this follows from the following observation: if $q(x)$ is any non-zero polynomial of degree strictly smaller than $p(x)$, then q is strictly positive or negative on some interval $[x_2, x_3]$ with $x_2, x_3 \in K_0$ and $x_0 < x_2 < \xi < x_3 < x_1$. So the sign of $q(\xi)$ in L_0 can be determined by checking the sign of $q(x_2)$ and the sign of $q(\xi')$ in L_1 can be determined by checking to sign of $q(f(x_2))$. But both answers should agree because f is an order preserving isomorphism.

So we have an isomorphism between L_0 and L_1. This isomorphism is necessarily unique because it should send the nth root from the left of the polynomial $p(x) \in K[x]$ in L_0 to the nth root from the left of $p(x)$ in L_1. □

6.5. Quantifier elimination.

THEOREM 2.41. *The theory RCOF of real closed ordered fields has quantifier elimination.*

PROOF. We use Theorem 1.5. So let K, L be two real closed ordered fields, where L in addition is ω_1-saturated, and suppose $f\colon \{k_1, \ldots, k_n\} \to L$ is a local isomorphism and $k \in K$. Then $\mathbb{Q}(k_1, \ldots, k_n)$, considered as an ordered subfield of K, and $\mathbb{Q}(f(k_1), \ldots, f(k_n))$, considered as an ordered subfield of L, are isomorphic. So we can use the previous theorem to extend f to an isomorphism $\overline{f}$ of ordered fields between the real closure $\overline{K}$ of $\mathbb{Q}(k_1, \ldots, k_n)$ inside K and the real closure $\overline{L}$ of $\mathbb{Q}(f(k_1), \ldots, f(k_n))$ inside L. If $k \in \overline{K}$, then we send k to $\overline{f}(k)$. So the interesting case is where k is transcendental over $\overline{K}$. To simplify notation, we will assume $\overline{K} = \overline{L}$.

In that case we should send k to an element $l \in L$ which is transcendental over the subfield $\overline{K}$ and for which

$$(\forall x \in \overline{K})\, x \le k \Leftrightarrow x \le l$$

holds. Such an element certainly exists because $|\overline{K}| = \omega$ and L is assumed to be ω^+-saturated. And this is enough, for to see that the composite isomorphism

$$\overline{K}(k) \cong \overline{K}(x) \cong \overline{K}(l)$$

is order preserving it suffices to check that $p(k)$ and $p(l)$ have the same sign for every irreducible polynomial $p \in \overline{K}[x]$. This is true for irreducible polynomials of degree one (by construction), and if p has degree greater than one, then p has no roots in K or L (since $\overline{K}$ is maximal as an algebraic extension over $\mathbb{Q}(k_1, \ldots, k_n)$ inside K or L). So p does not change sign inside K or L and $p(k)$ and $p(l)$ have the same sign as $p(0)$. □

COROLLARY 2.42. *The theory RCOF is complete.*

PROOF. Since the theory of real closed ordered fields has quantifier elimination and has a model which can be embedded into any other model (to wit, the real numbers which are algebraic over $\mathbb{Q}$), this theory is complete by Theorem 1.7. □

REMARK 2.43. The theory RCOF is not λ-categorical for any infinite λ, but that is not so easy to prove!

6.6. Hilbert's 17th Problem.

THEOREM 2.44. (Hilbert's 17th Problem) *Let K be a real closed field. If $f \in K(x_1, \ldots, x_n)$ is such that $f(a_1, \ldots, a_n) \geq 0$ for all $a_1, \ldots, a_n \in K$, then f can be written as*

$$f = g_1^2 + \ldots + g_n^2$$

for suitable $g_i \in K(x_1, \ldots, x_n)$.

PROOF. Suppose f cannot be written as a sum of squares in $K(x_1, \ldots, x_n)$. The same applies to -1, because -1 cannot be written as a sum of squares in K. So we can order $K(x_1, \ldots, x_n)$ in such a way that f becomes negative. This order extends the original order on K because K is real closed and hence positive elements in K can be written as squares (see Proposition 2.35). Now embed $K(x_1, \ldots, x_n)$ with this order into a real closed field L. So we have embeddings of fields

$$K \subseteq K(x_1, \ldots, x_n) \subseteq L,$$

all of which preserve and reflect the ordering. So the inclusion $K \subseteq L$ reflects truth of atomic sentences, and hence of quantifier-free sentences and hence, as the theory of real closed fields has quantifier elimination, of *all* sentences. Therefore the sentence

$$\exists x_1, \ldots, x_n \, f(x_1, \ldots, x_n) < 0,$$

which is true in L, must be true in K as well. □

REMARK 2.45. Hilbert's 17th Problem asked whether Theorem 2.44 holds in case K is the reals. This was settled by Artin in 1927, who proved the result for general real closed fields. The model-theoretic proof we just gave is due to Robinson.