Model Theory

Course Notes

May 4, 2025

Manuel Bodirsky, Institut für Algebra, TU Dresden

Disclaimer: this is a draft and probably contains many typos and mistakes. Please report them to Manuel.Bodirsky@tu-dresden.de.

The text contains 100 exercises; the ones with a star are harder.

Acknowledgements. The author is grateful to Leopold Schlicht, Stephanie Feilitzsch, Florian Starke, and to the students at TU Dresden from the summer semester 2021 for reporting typos and mistakes.

Contents

CHAPTER 1

First-order Logic

There are many excellent text books about model theory. A classic is Hodges' *model theory* [**18**]; a short version is also available [**19**]. Another introduction is given by Marker [**22**]. A more recent book has been written by Tent and Ziegler [**29**].

1.1. First-order Structures

1.1.1. Signatures. A *signature* τ is a set of relation and function symbols, each equipped with an *arity* $k \in \mathbb{N}$. Examples of important signatures:

- $\tau_{\text{Graph}} = \{E\}$
- $\tau_{\text{AGroup}} = \{+, -, 0\}$
- $\tau_{\text{Ring}} = \tau_{\text{AG}} \cup \{1, \cdot\}$
- $\tau_{\text{Group}} = \{\circ, ^{-1}, e\}$
- $\tau_{\text{LO}} = \{<\}$
- $\tau_{\text{ORing}} = \tau_{\text{Ring}} \cup \tau_{\text{LO}}$
- $\tau_{\text{Arithmetic}} = \{0, s, +, \cdot, <\}$
- $\tau_{\text{Set}} = \{\in\}$

Here, $+$, $\circ$, and $\cdot$ are binary function symbols, $-$, $^{-1}$, and s are unary function symbols, e, 0, and 1 are 0-ary function symbols, and E, $\in$, and $<$ are binary relation symbols.

1.1.2. Structures. A *τ-structure* $\underline{A}$ is a set A (called the *domain*, or the *base set*, or the *universe* of $\underline{A}$) together with

- a relation $R^{\underline{A}} \subseteq A^k$ for each k-ary relation symbol $R \in \tau$. Here we allow the case $k = 0$, in which case $R^{\underline{A}}$ is either empty or of the form $\{()\}$, i.e, the set consisting of the empty tuple;
- a function $f^{\underline{A}} \colon A^k \to A$ for each k-ary function symbol $f \in \tau$; here we also allow the case $k = 0$ to model constants from A.

Unless stated otherwise, $A, B, C, \dots$ denotes the domain of the structure $\underline{A}$, $\underline{B}$, $\underline{C}$, $\dots$, respectively. We sometimes write $(A; R_1^{\underline{A}}, R_2^{\underline{A}}, \dots, f_1^{\underline{A}}, f_2^{\underline{A}}, \dots)$ for the structure $\underline{A}$ with relations $R_1^{\underline{A}}, R_2^{\underline{A}}, \dots$ and functions $f_1^{\underline{A}}, f_2^{\underline{A}}, \dots$ Well-known examples of structures are $(\mathbb{Q}; +, -, 0, 1, \cdot, <)$, $(\mathbb{C}; +, -, 0, 1, \cdot, <)$, $(\mathbb{Z}; 0, s, +, \cdot, <)$, etc. We say that a structure is infinite if its domain is infinite. If all symbols in τ are function symbols, then a τ-structure $\underline{A}$ is also called an *algebra* (in the sense of universal algebra); in the other extreme, if all symbols in τ are relation symbols, then $\underline{A}$ is called a *relational structure*.

EXAMPLE 1. A *(simple, undirected) graph* is a pair (V, E) consisting of a set of *vertices* V and a set of *edges* $E \subseteq \binom{V}{2}$, that is, E is a set of 2-element subsets of V. Graphs can be modelled as relational structures $\underline{G}$ using a signature that contains a single binary relation symbol R, putting $G := V$ and adding (u, v) to $R^{\underline{G}}$ if $\{u, v\} \in E$. If we insist that a structure with this signature satisfies $(x, y) \in R^{\underline{G}} \Rightarrow (y, x) \in R^{\underline{G}}$ and $(x, x) \notin R^{\underline{G}}$, then we can associate to such a structure an undirected graph and obtain a bijective correspondence between undirected graphs with vertices V and structures $\underline{G}$ with domain V as described above. △

EXAMPLE 2. A *group* is an algebra $\underline{G}$ with the signature $\tau_{\text{Group}} = \{\circ, ^{-1}, e\}$ such that for all $x, y, z \in G$ we have that

- $x \circ (y \circ z) = (x \circ y) \circ z$,
- $x \circ x^{-1} = e$,
- $e \circ x = x$ and $x \circ e = x$. △

1.1.3. Substructures. A τ-structure $\underline{A}$ is a *substructure* of a τ-structure $\underline{B}$ iff

- $A \subseteq B$,
- for each $R \in \tau$, and for all tuples $\bar{a}$ from A, $\bar{a} \in R^{\underline{A}}$ iff $\bar{a} \in R^{\underline{B}}$, and
- for each $f \in \tau$ we have that $f^{\underline{A}}(\bar{a}) = f^{\underline{B}}(\bar{a})$.

In this case, we also say that $\underline{B}$ is an *extension* of $\underline{A}$. Substructures $\underline{A}$ of $\underline{B}$ and extensions $\underline{B}$ of $\underline{A}$ are called *proper* if the domains of $\underline{A}$ and $\underline{B}$ are distinct.

Note that for every subset S of the domain of $\underline{B}$ there is a unique smallest substructure of $\underline{B}$ whose domain contains S, which is called the *substructure of $\underline{B}$ generated by S*, and which is denoted by $\underline{B}[S]$.

EXAMPLE 3. When we view a graph as an $\{E\}$-structure $\underline{G}$, then a subgraph is not necessarily a substructure of $\underline{G}$. In graph theory, the substructures of $\underline{G}$ are called *induced subgraphs*: the difference is that in an induced subgraph (V', E') of (V, E) the edge set must be of the form $E' := E \cap \binom{V'}{2}$ instead of an arbitrary subset of it. △

EXAMPLE 4. Due to the choice of our signature τ_{Group}, the subgroups of $\underline{G}$ are precisely the substructures of $\underline{G}$ as defined above. △

1.1.4. Expansions and Reducts. Let σ, τ be signatures with $\sigma \subseteq \tau$. If $\underline{A}$ is a σ-structure and $\underline{B}$ is a τ-structure, both with the same domain, such that $R^{\underline{A}} = R^{\underline{B}}$ for all relations $R \in \sigma$ and $f^{\underline{A}} = f^{\underline{B}}$ for all functions and constants $f \in \sigma$, then $\underline{A}$ is called a *reduct* of $\underline{B}$, and $\underline{B}$ is called an *expansion* of $\underline{A}$.

Let $\underline{A}$ be a τ-structure and $B \subseteq A$. We write $\underline{A}_B$ for the expansion of $\underline{A}$ by the constant symbol b for every element $b \in B$ (we set $b^{\underline{A}_B} := b$).

1.1.5. Homomorphisms. In the following, let $\underline{A}$ and $\underline{B}$ be τ-structures. A *homomorphism* h from $\underline{A}$ to $\underline{B}$ is a mapping from A to B that *preserves* each function and each relation for the symbols in τ; that is,

- if $(a_1, \dots, a_k)$ is in $R^{\underline{A}}$, then $(h(a_1), \dots, h(a_k))$ must be in $R^{\underline{B}}$;
- $f^{\underline{B}}(h(a_1), \dots, h(a_k)) = h(f^{\underline{A}}(a_1, \dots, a_k))$.

A homomorphism from $\underline{A}$ to $\underline{B}$ is called a *strong homomorphism* if it also preserves the complements of the relations from $\underline{A}$. Injective strong homomorphisms are called *embeddings*. Surjective embeddings are called *isomorphisms*.

EXAMPLE 5. Group homomorphisms, field endomorphisms, ring endomorphisms, linear maps between vector spaces. △

EXAMPLE 6. The *graph colorability problem* is an important problem in discrete mathematics, with many applications in theoretical computer science (it can be used to model e.g. frequency assignment problems). As a computational problem, the graph n-colorability problem has the following form.

Given: a finite graph $G = (V, E)$.
Question: can we colour the vertices of G with n colours such that adjacent vertices get different colours?

The n-colorability problem can be formulated as a graph homomorphism problem: is there a homomorphism from G to

$$K_n := (\{1, \dots, n\}; E^{K_n}) \text{ where } E^{K_n} := \{(u, v) \in V^2 \mid u \neq v\}\,.$$

We also refer to these homomorphisms as *proper colourings* of G, and say that G is n-*colourable* if such a colouring exists. The *chromatic number* $\chi(G)$ of G is the minimal natural number $n \in \mathbb{N}$ such that G is n-colourable. For example, the chromatic number of K_n is n. △

EXAMPLE 7. We present a concrete instance of a colouring problem from pure mathematics. Let (V, E) the unit distance graph on $\mathbb{R}^2$, i.e., the graph has the vertex set $V := \mathbb{R}^2$ (we imagine the nodes as the points of the Euclidean plane) and edge set

$$E := \big\{(x, y) \in V^2 \;\big|\; |x - y| = 1\big\}\,.$$

In other words, two points are linked by an edge if they have distance one. What is the chromatic number of this graph? △

The following statement follows from the compactness theorem of first-order logic, as we will see later in this chapter.

PROPOSITION 1.1.1. *Let G be a graph such that all finite subgraphs of G are k-colourable. Then G is k-colourable.*

The problem to determine the chromatic number χ of this graph is known as the Hadwiger-Nelson problem. It is known that $\chi \leq 7$. We have seen that $4 \leq \chi$. In April 11, 2018, Aubrey de Grey announced a proof that $5 \leq \chi$. The precise value of $\chi \in \{5, 6, 7\}$ is not known.

1.1.6. Products. Let $\underline{A}$ and $\underline{B}$ be τ-structures. Then the (*direct*, or *categorical*) *product* $\underline{C} = \underline{A} \times \underline{B}$ is the τ-structure with domain $A \times B$, which has for each k-ary $R \in \tau$ the relation that contains a tuple $((a_1, b_1), \dots, (a_k, b_k))$ if and only if $R(a_1, \dots, a_k)$ holds in $\underline{A}$ and $R(b_1, \dots, b_k)$ holds in $\underline{B}$. For each k-ary $f \in \tau$ the structure $\underline{C}$ has the operation that maps $((a_1, b_1), \dots, (a_k, b_k))$ to $(f(a_1, \dots, a_k), f(b_1, \dots, b_k))$. The direct product $\underline{A} \times \underline{A}$ is also denoted by $\underline{A}^2$, and the k-fold product $\underline{A} \times \cdots \times \underline{A}$, defined analogously, by $\underline{A}^k$.

Note that the product of two groups, viewed as structures, is precisely the well-known product of groups.

We generalise the definition of products to infinite products as follows. For a sequence of τ-structures $(\underline{A}_i)_{i \in I}$, the direct product $\underline{B} = \prod_{i \in I} \underline{A}_i$ is the τ-structure on the domain $\prod_{i \in I} A_i$ such that for $R \in \tau$ of arity k

$$\big((a_i^1)_{i \in I}, \dots, (a_i^k)_{i \in I}\big) \in R^{\underline{B}} \text{ iff } (a_i^1, \dots, a_i^k) \in R^{\underline{A}_i} \text{ for each } i \in I\;,$$

and for $f \in \tau$ of arity k, we have

$$f^{\underline{B}}\big((a_i^1)_{i \in I}, \dots, (a_i^k)_{i \in I}\big) = \big(f^{\underline{A}_i}(a_i^1, \dots, a_i^k)\big)_{i \in I}\,.$$

If $\mathbf{A}_i = \mathbf{A}$ for all $i \in I$, we also write $\mathbf{A}^I$ instead of $\prod_{i \in I} \underline{A}_i$, and call it the *$I$-th power of $\underline{A}$*.

Exercises.

(1) Show that there is a homomorphism from $\underline{A}^2$ to $\underline{A}$. Find an example of a structure $\underline{A}$ where all such homomorphisms are strong, and another example where all such homomorphisms are not strong.
(2) Find a signature τ for vector spaces and describe how a vector space may be viewed as a τ-structure. Your definition should have the property that homomorphisms between the structures you consider correspond precisely to linear maps.

(3) Let $\underline{A}$ and $\underline{B}$ be τ-structures and suppose that $G \subseteq A$ generates $\underline{A}$, i.e., $\underline{A} = \underline{A}[G]$. Then every homomorphism $h\colon \underline{A} \to \underline{B}$ is determined by its values on G.

(4) Let $\underline{A}$ and $\underline{B}$ be τ-structures, and suppose that τ has no relation symbols. Show that every bijective homomorphism from $\underline{A}$ to $\underline{B}$ is an isomorphism.

(5) Find an example of a τ-structure $\underline{A}$ and a bijective homomorphism from $\underline{A}$ to $\underline{A}$ which is not an isomorphism.

1.2. Formulas, Sentences, Theories

To define the syntax of first-order logic, we first have to introduce terms, before we can define (first-order) formulas and (first-order) sentences, and finally (first-order) theories.

1.2.1. Terms. Let τ be a signature. In this section we will see how to use the function symbols in τ to build *terms.*

Definition 1.2.1. A *(τ-)term* is defined inductively:

- constants from τ are τ-terms;
- variables $x_0, x_1, \ldots$ are τ-terms;
- if $t_1, \ldots, t_k$ are τ-terms, and $f \in \tau$ has arity k, then $f(t_1, \ldots, t_k)$ is a τ-term.

We write $t(x_1, \ldots, x_n)$ if all variables that appear in t come from $\{x_1, \ldots, x_n\}$; we do not require that each variable x_i appears in t.

Example 8. Well-known examples of terms are polynomials over a ring $\underline{R}$: they are terms over a signature that contains a binary symbol $+$ for addition and a binary symbol $*$ for multiplication, together with a constant symbol for each element of R. △

1.2.2. Semantics of Terms. In this section we describe how τ-terms over a given τ-structure describe functions (in the same way as polynomials describe polynomial functions over a given ring).

Let $\underline{A}$ be a τ-structure and let $x_1, \ldots, x_n$ be distinct variables. Every τ-term $t(x_1, \ldots, x_n)$ describes a function $t^{\underline{A}}\colon A^n \to A$ as follows:

(1) if t equals $c \in \tau$ then $t^{\underline{A}}$ is the function $(a_1, \ldots, a_n) \mapsto c^{\underline{A}}$;
(2) if t equals x_i then $t^{\underline{A}}$ is the function $(a_1, \ldots, a_n) \mapsto a_i$;
(3) if t equals $f(t_1, \ldots, t_k)$ for a k-ary $f \in \tau$ then $t^{\underline{A}}$ is the function

$$(a_1, \ldots, a_n) \mapsto f^{\underline{A}}\big(t_1^{\underline{A}}(a_1, \ldots, a_n), \ldots, t_k^{\underline{A}}(a_1, \ldots, a_n)\big)$$

Note that item 1 is a special case of item 3. The function described by t is also called the *term function* of t (with respect to $\underline{A}$).

Exercises.

(6) Let $\underline{B}$ be a τ-structure and $G \subseteq B$. Let $\underline{A}$ be the substructure of $\underline{B}$ generated by G. Show that for every element $a \in A$ there exists a τ-term $t(x_1, \ldots, x_n)$ and elements $g_1, \ldots, g_n \in G$ such that $t^{\underline{B}}(g_1, \ldots, g_n) = a$.

1.2.3. Formulas and sentences. Let τ be a signature. The relation symbols in the signature τ did not play any role when defining τ-terms, but they become important when defining *τ-formulas.* Moreover, the equality symbol $=$ is 'hard-wired' into first-order logic; we can use it to create formulas by equating terms. Finally, we can combine formulas using Boolean connectives, and quantify over variables.

Definition 1.2.2. An *atomic (τ-)formula* is an expression of the form

- $t_1 = t_2$ where t_1 and t_2 are τ-terms;

- $R(t_1,\dots,t_k)$ where $t_1,\dots,t_k$ are τ-terms and $R \in \tau$ is a k-ary relation symbol.

Formulas are defined inductively as follows:

- atomic formulas are formulas;
- if ϕ is a formula, then $\neg\phi$ is a formula *(negation)*;
- if ϕ and ψ are formulas, then $\phi \wedge \psi$ is a formula *(conjunction)*;
- if ϕ is a formula, and x is a variable, then $\exists x.\, \phi$ is a formula *(existential quantification)*

Atomic formulas and negations of atomic formulas are sometimes called *literals*. Similarly as for terms, we write $\phi(x_1,\dots,x_n)$ if all non-quantified variables in the formula ϕ come from $\{x_1,\dots,x_n\}$. A *(first-order) sentence* is a formula without free variables, i.e., all variables that appear in the formula are *bound*, i.e., quantified by some quantifier.

EXAMPLE 9. Let $\tau = \{R\}$ where R is a binary relation symbol. Then the following formula is an example of a first-order sentence

$$\forall x_1, x_2, x_3 \big(R(x_1,x_2) \wedge R(x_2,x_3) \Rightarrow R(x_1,x_3)\big). \qquad \triangle$$

1.2.4. Semantics of formulas. So far, we have just introduced the *syntax* of first-order logic, i.e., the shape of formulas and sentences, without discussing what these expressions actually mean. In this section we discuss their *semantics*; the idea is that formulas can be used to define new relations and functions in a given structure, and that sentences can be used to describe properties that a structure might have or not.

Let $\underline{A}$ is a τ-structure. Then every τ-formula $\phi(x_1,\dots,x_n)$ describes a relation $\phi^{\underline{A}} \subseteq A^n$ as follows:

- if ϕ equals $t_1 = t_2$ then $\phi^{\underline{A}}$ is the relation
$$\big\{(a_1,\dots,a_n) \mid t_1^{\underline{A}}(a_1,\dots,a_n) = t_2^{\underline{A}}(a_1,\dots,a_n)\big\};$$
- if ϕ equals $R(t_1,\dots,t_k)$ then
$$\phi^{\underline{A}} := \big\{(a_1,\dots,a_n) \mid (t_1^{\underline{A}}(\bar a),\dots,t_k^{\underline{A}}(\bar a)) \in R^{\underline{A}}\big\};$$
- if ϕ equals $\phi_1 \wedge \phi_2$ then $\phi^{\underline{A}} := \phi_1^{\underline{A}} \cap \phi_2^{\underline{A}}$;
- if ϕ equals $\neg\psi$ then $\phi^{\underline{A}} := A^n \setminus \psi^{\underline{A}}$;
- if ϕ equals $\exists x.\, \psi(x, x_1,\dots,x_n)$ then
$$\phi^{\underline{A}} := \{(a_1,\dots,a_n) \mid \text{there exists } a \in A \text{ s.t. } (a, a_1,\dots,a_n) \in \psi^{\underline{A}}\}.$$

A relation $R \subseteq A^k$ is called *(first-order) definable in* $\underline{A}$ if there exists a τ-formula $\phi(x_1,\dots,x_k)$ such that $R = \phi^{\mathbf{A}}$; we also say that ϕ *defines* R over $\underline{A}$. An function $f\colon S \to T$, for $S \subseteq A^k$ and $T \subseteq A^l$, is called definable if the relation

$$\{(a_1,\dots,a_k,b_1,\dots,b_l) \mid f(a_1,\dots,a_k) = (b_1,\dots,b_l)\}$$

is definable in $\underline{A}$.

We freely use brackets to avoid ambiguities when writing down terms. That is, we write $\neg(x \wedge y)$ to distinguish it from $(\neg x) \wedge y$. If brackets are omitted, there is the convention that negation $\neg$ binds stronger than conjunction $\wedge$; that is, $\neg x \wedge y$ would stand for $(\neg x) \wedge y$.

For $\phi(x_1,\dots,x_n)$ we write

$$\underline{A} \models \phi(a_1,\dots,a_n)$$

instead of $(a_1,\dots,a_n) \in \phi^{\underline{A}}$. In particular, if ϕ is a sentence, i.e., if $n = 0$, we write $\underline{A} \models \phi$, and say that $\underline{A}$ *satisfies* ϕ, if $() \in \phi^{\underline{A}}$ (that is, if $\phi^{\underline{A}} \neq \emptyset$).

Shortcuts:

- *Disjunction:* $\phi \vee \psi$ is an abbreviation for $\neg(\neg\phi \wedge \neg\psi)$
- *Implication:* $\phi \Rightarrow \psi$ is an abbreviation for $\neg\phi \vee \psi$
- *Equivalence:* $\phi \Leftrightarrow \psi$ is an abbreviation for $\phi \Rightarrow \psi \wedge \psi \Rightarrow \phi$
- *Universal quantification:* $\forall x.\, \phi(x)$ is an abbreviation for $\neg\exists x.\, \neg\phi(x)$
- *Inequality:* $x \neq y$ is an abbreviation for $\neg(x = y)$
- *False:* $\perp$ is an abbreviation for $\exists x. x \neq x$.
- *True:* $\top$ is an abbreviation for $\neg\perp$ (the same as $\forall x. x = x$).

Moreover, when A is a unary relation symbol, then we may write $\exists x \in A.\, \phi$ instead of $\exists x\, (A(x) \wedge \phi)$ and $\forall x \in A.\, \phi$ instead of $\forall x\, (A(x) \Rightarrow \phi)$.

Example 10. The following statements about well-known structures follow straightforwardly from the definitions.

- $(\mathbb{Z}; <) \models 0 < 1$
- $(\mathbb{Q}; <) \models \forall x, y \big(x < y \Rightarrow \exists z\, (x < z \wedge z < y)\big)$ (density)
- $(\mathbb{Z}; <) \not\models \forall x, y \big(x < y \Rightarrow \exists z\, (x < z \wedge z < y)\big)$ △

1.2.5. First-order Theories. Let τ be a signature. A *(τ-)theory* is a set of first-order τ-sentences. Let $\underline{A}$ be a τ-structure and T a τ-theory. Then $\underline{A}$ is a *model of* T, in symbols $\underline{A} \models T$, if $\underline{A} \models \phi$ for all $\phi \in T$. A τ-theory T is called

- *satisfiable* (or *consistent*) if T has a model.
- *complete* if T is satisfiable and for every τ-sentence either $\phi \in T$ or $\neg\phi \in T$.

The *(first-order) theory of* $\underline{A}$ is defined as the set of all first-order τ-sentences that are satisfied by $\underline{A}$. Note that $\mathrm{Th}(\underline{A})$ is always a complete theory.

Example 11. A famous example of a first-order theory is the set ZF of axioms of Zermelo-Fraenkl set theory. Recall that this is an infinite set of first-order sentences over the signature $\tau = \{\in\}$. An example of a sentence in ZF is

$$\exists x \forall y.\, \neg(y \in x)$$

stating the existence of the empty set. Note that ZF is not complete: for example the Axiom of Choice (C) is independent from ZF in the sense that both $\mathrm{ZFC} := \mathrm{ZF} \cup \{\mathrm{C}\}$ and $\mathrm{ZF} \cup \{\neg\mathrm{C}\}$ have a model. △

Example 12. The theory T_{AGroup} of abelian groups is over the signature $\tau_{\mathrm{AGroup}} = \{+, -, 0\}$ and contains the following axioms:

$$\begin{array}{ll}
\forall x, y, z.\ (x + y) + z = x + (y + z) & \text{(associativity)} \\
\forall x.\ 0 + x = x + 0 = y & \text{(neutral element)} \\
\forall x.\ x + (-x) = 0 & \text{(inverse elements)} \\
\forall x, y.\ x + y = y + x & \text{(Abelian)}
\end{array}$$
△

Example 13. The theory T_{CRing} of commutative rings is over the signature τ_{Ring}, contains T_{AGroup} and the following additional axioms:

$$\begin{array}{ll}
\forall x, y, z.\ (xy)z = x(yz) & \text{(associativity)} \\
\forall x.\ 1 \cdot x = x & \text{(multiplicative unit)} \\
\forall x, y.\ xy = yx & \text{(commutativity)} \\
\forall x, y, z.\ x(y + z) = xy + xz & \text{(distributivity)}
\end{array}$$
△

EXAMPLE 14. The theory T_{Field} of fields is over the signature τ_{Ring}, contains T_{CRing} and the following additional axioms:

$$\neg(0 = 1)$$
$$\forall x\ (\neg(x = 0) \Rightarrow \exists y.xy = 1) \qquad \triangle$$

If S and T are τ-theories then we write

$$S \models T$$

if every model of S is also a model of T. We also write $S \models \phi$ instead of $S \models \{\phi\}$.

Exercises.

(7) Let $\underline{A}$ and $\underline{B}$ be τ-structures. If $a \in A^n$ and $b \in B^n$ satisfy the same quantifier-free formulas, then the substructure of $\underline{A}$ generated by the entries of a is isomorphic to the substructure of $\underline{B}$ generated by the entries of b.

(8) Show that for every first-order formula $\phi(x)$ and for every $\ell \in \mathbb{N}$ there exists a first-order sentence ψ such that a structure $\underline{A}$ satisfies ψ if and only if there are precisely ℓ elements $a \in A$ such that $\underline{A} \models \phi(a)$. The sentence ψ is often denoted by $\exists x^{=\ell}.\phi$.

(9) Write down the axioms of algebraically closed fields in first-order logic.

(10) A formula is in *prenex normal form* if it is of the form $Q_1 x_1 \dots Q_n x_n.\, \phi$ where Q_i is either $\exists$ or $\forall$ and ϕ is without quantifiers. Show that every formula $\phi(y_1, \dots, y_n)$ is *equivalent* to a formula $\psi(y_1, \dots, y_n)$ in prenex normal form, that is, for every structure $\underline{M}$ we have $\underline{M} \models \forall \bar{y}(\phi(\bar{y}) \Leftrightarrow \psi(\bar{y}))$.

(11) Show that every quantifier-free formula is equivalent to a formula in *conjunctive normal form*, i.e., to a conjunction of disjunctions of atomic formulas and negated atomic formulas.

1.3. Vaught's Conjecture

Let T be a complete theory over a countable signature and let κ be a cardinal. How many models of cardinality κ can T have, up to isomorphism? Let $I_T(\kappa)$ be the number of isomorphism types of models of T of cardinality κ, also called the *spectrum* of T. Examples:

$$I_{\text{Th}(K_5)}(\aleph_0) = 0 \qquad \text{(See Exercise 12)}$$
$$I_{\text{Th}(\mathbb{Q};<)}(\aleph_0) = 1 \qquad \text{(see Example 15)}$$
$$I_{\text{Th}(\mathbb{Z};\{(x,y) \mid y=x+1\})}(\aleph_0) = \aleph_0 \qquad \text{(exercise)}$$
$$I_{\text{Th}(\mathbb{R};+,<)}(\aleph_0) = 2^{\aleph_0} \qquad \text{(exercise)}$$

EXAMPLE 15. $I_{\text{Th}(\mathbb{Q};<)}(\aleph_0) = 1$. The standard approach to verify this is a so-called *back-and-forth argument*. Let $\underline{A}$ be a countable model of the first-order theory T of $(\mathbb{Q};<)$. It is easy to verify that T contains (and, as this argument will show, is uniquely determined by)

- $\exists x.\ x = x$ (no empty model)
- $\forall x, y, z\ ((x < y \wedge y < z) \Rightarrow x < z)$ (transitivity)
- $\forall x.\ \neg(x < x)$ (irreflexivity)
- $\forall x, y\ (x < y \vee y < x \vee x = y)$ (totality)
- $\forall x\, \exists y.\ x < y$ (no largest element)
- $\forall x\, \exists y.\ y < x$ (no smallest element)
- $\forall x, z\ \big(x < z \Rightarrow \exists y\ (x < y \wedge y < z)\big)$ (density).

An isomorphism between $\underline{A}$ and $(\mathbb{Q};<)$ can be defined inductively as follows. Suppose that we have already defined f on a finite subset S of $\mathbb{Q}$ and that f is an embedding of the structure induced by S in $(\mathbb{Q};<)$ into $\underline{A}$. Since $<^{\underline{A}}$ is dense and unbounded, we can extend f to any other element of $\mathbb{Q}$ such that the extension is still an embedding from a substructure of $\mathbb{Q}$ into $\underline{A}$ (*going forth*). Symmetrically, for every element v of $\underline{A}$ we can find an element $u \in \mathbb{Q}$ such that the extension of f that maps u to v is also an embedding (*going back*). We now alternate between going forth and going back; when going forth, we extend the domain of f by the *next* element of $\mathbb{Q}$, according to some fixed enumeration of the elements in $\mathbb{Q}$. When going back, we extend f such that the image of A contains the *next* element of $\underline{A}$, according to some fixed enumeration of the elements of $\underline{A}$. If we continue in this way, we have defined the value of f on all elements of $\mathbb{Q}$. Moreover, f will be surjective, and an embedding, and hence an isomorphism between $\underline{A}$ and $(\mathbb{Q};<)$. $\triangle$

Vaught's theorem states that for any theory T the spectrum $I_T(\aleph_0)$ cannot be two. Morley showed that if $I_T(\aleph_0)$ is infinite, then

$$I_T(\aleph_0) \in \{\aleph_0, \aleph_1, 2^{\aleph_0}\}$$

Vaught's Conjecture: if $I_T(\aleph_0)$ is infinite, then $I_T(\aleph_0) \in \{\aleph_0, 2^{\aleph_0}\}$.

Exercises.

(12) Show that two finite structures are isomorphic if and only if they have the same theory.

(13) Show that if τ is a countable signature, then for any infinite cardinal κ there are at most 2^κ many non-isomorphic τ-structures of cardinality κ.

(14) Let $\underline{A}$ be a $\{E\}$-structure where E is an equivalence relation on A with infinitely many infinite classes. Show that $I_{\mathrm{Th}(\underline{A})}(\aleph_0) = 1$.

(15) $(*)$ Prove the claim that $I_{\mathrm{Th}(\mathbb{Z};\{(x,y) \mid y=x+1\})}(\aleph_0) = \aleph_0$.

(16) $(*)$ Determine the value of $I_{\mathrm{Th}(\mathbb{Z};<)}(\aleph_0)$.

CHAPTER 2

Compactness

A special case of the compactness theorem was already proved by Gödel; the general case below was proved by A. Maltsev in 1936.

THEOREM 2.0.1. *A first-order theory T is satisfiable if and only if all finite subsets of T are satisfiable.*

REMARK 2.0.2. The name *compactness theorem* comes from the fact that the compactness theorem is equivalent to the statement that the following natural topological space is compact: the space is the set $\mathcal{T}(\tau)$ of all complete τ-theories, and the basic open sets are the sets T_ϕ of the form $\{T \in \mathcal{T}(\tau) \mid \phi \in T\}$.

To see the equivalence, let $\mathcal{C}$ be a covering of $\mathcal{T}(\tau)$ by open subsets of $\mathcal{T}(\tau)$. This means that $\mathcal{T}(\tau)$ can be written as $\bigcup_{\phi \in S} T_\phi$ for some set S of τ-sentences. Consider the set S' of all τ-sentences that are not contained in S. If S' has a model $\underline{B}$, then the theory $\mathrm{Th}(\underline{B})$ is not covered by $\mathcal{C}$. The compactness theorem implies that there is a finite subset F of S' which is unsatisfiable. But then $\{T_{\neg\phi} \mid \phi \in F\}$ is a finite subset of $\mathcal{C}$ covering $\mathcal{T}(\tau)$, showing that $\mathcal{T}(\tau)$ is compact.

Conversely, suppose that $\mathcal{T}(\tau)$ is compact, and that the τ-theory S is unsatisfiable. Then $\{T_{\neg\phi} \mid \phi \in S\}$ is an open covering of $\mathcal{T}(\tau)$. So by compactness it has a finite subcovering, i.e., there is a finite subset F of S such that $\bigcup\{T_{\neg\phi} \mid \phi \in F\} = \mathcal{T}(\tau)$. Hence, F is unsatisfiable, which is the statement of the compactness theorem. $\square$

There are several different proofs of the compactness theorem. It is a consequence of the completeness theorem of first-order logic, and one can show it using Henkin's construction. In this course we present a proof based on *ultraproducts*; they are an important tool of model theory to build interesting structures and have found many applications in mathematics, e.g., in topology, set theory, and algebra.

Exercises.

(17) Show that if a first-order theory has infinite models, then it also has arbitrarily large models (for every set X there is a model $\underline{M}$ with $|M| \geq |X|$).

(18) Let $\underline{K}$ be a field. The *characteristic* of $\underline{K}$ is the smallest $n \in \mathbb{N}$ such that

$$\underbrace{1 + \cdots + 1}_{n \text{ times}} = 0$$

if such an n exists, and 0 otherwise. Let T be a theory that contains the theory of fields.

- Show that if there are models of T of arbitrarily large finite characteristic, then there exists a model of T of characteristic 0.
- Show that there is no finite first-order theory whose models are precisely the fields of characteristic 0.

2.1. Filter

Let X be a set. A filter on X is a certain set of subsets of X; the idea is that the elements of $\mathcal{F}$ are (in some sense) 'large'; it helps thinking of the elements $F \in \mathcal{F}$ as being 'almost all' of X.

DEFINITION 2.1.1. A *filter* $\mathcal{F}$ on X is a set of subsets of X such that

(1) $\emptyset \notin \mathcal{F}$ and $X \in \mathcal{F}$;
(2) if $F \in \mathcal{F}$ and $G \subseteq X$ contains F, then $G \in \mathcal{F}$.
(3) if $F_1, F_2 \in \mathcal{F}$ then $F_1 \cap F_2 \in \mathcal{F}$.

Note that filters have the *finite intersection property*:

$$A_1, \dots, A_n \in \mathcal{F} \Rightarrow A_1 \cap \dots \cap A_n \neq \emptyset \qquad \text{(FIP)}$$

LEMMA 2.1.2. *Every subset $\mathcal{S} \subseteq \mathcal{P}(X)$ with the FIP is contained in a smallest filter that contains $\mathcal{S}$; this filter is called the filter* generated *by $\mathcal{S}$.*

PROOF. First add finite intersections, and then all supersets to $\mathcal{S}$. □

EXAMPLE 16. For a non-empty subset $Y \subseteq X$, the family

$$\mathcal{F} := \{Z \subseteq X \mid Y \subseteq Z\}$$

is a filter, the filter generated by a $\{Y\}$; such filters are called *principal*. △

EXAMPLE 17. The *Fréchet filter*: for an infinite set X this is the filter $\mathcal{F}$ that consists of all cofinite subsets of X, i.e.,

$$\mathcal{F} := \{Y \subseteq X \mid X \setminus Y \text{ is finite}\}. \qquad \triangle$$

2.2. Ultrafilter

A filter $\mathcal{F}$ is called a *ultrafilter* if $\mathcal{F}$ is maximal, that is for every filter $\mathcal{G} \supseteq \mathcal{F}$ we have $\mathcal{G} = \mathcal{F}$.

LEMMA 2.2.1. *Let $\mathcal{F}$ be a filter. Then the following are equivalent.*

(1) $\mathcal{F}$ is a ultrafilter.
(2) For all $A \subseteq X$ either $A \in \mathcal{F}$ or $X \setminus A \in \mathcal{F}$.
(3) For all $A_1 \cup \dots \cup A_n \in \mathcal{F}$ there is an $i \leq n$ with $A_i \in \mathcal{F}$.

PROOF. (1) $\Leftarrow$ (2): No $A \subseteq X$ can be added to $\mathcal{F}$. Hence, $\mathcal{F}$ is maximal.
(2) $\Leftarrow$ (3): Note that $A \cup (X \setminus A) = X \in \mathcal{F}$.
(1) $\Rightarrow$ (3): If there is an $i \leq n$ such that $\mathcal{F} \cup \{A_i\}$ has the FIP, then by Lemma 2.1.2 there is a filter that contains this set, and hence $\mathcal{F}$ was not maximal. Otherwise, there are $S_1, \dots, S_n \subseteq \mathcal{F}$ with $A_i \cap S_i = \emptyset$. Then $S_i \subseteq X \setminus A_i$ and thus $S_1 \cap \dots \cap S_n \subseteq X \setminus (A_1 \cup \dots \cup A_n) \notin \mathcal{F}$, a contradiction. □

A filter $\mathcal{F}$ is *principal* if it contains a inclusionwise minimal element. Note that this is the case if and only if $\bigcap \mathcal{F} \in \mathcal{F}$.

LEMMA 2.2.2. *Let $\mathcal{F}$ be a filter on a set X. Then the following are equivalent.*

(1) $\mathcal{F}$ is a principal ultrafilter;
(2) $\mathcal{F}$ contains $\{a\}$ for some $a \in X$.
(3) $\mathcal{F}$ is of the form $\{Y \subseteq X \mid a \in Y\}$ for some $a \in X$.
(4) $\mathcal{F}$ is an ultrafilter and contains a finite set.

PROOF. (1) ⇒ (2): let $A := \bigcap \mathcal{F} \in \mathcal{F}$. If $|A| > 1$ then we can write $A = B_1 \cup B_2$ for $B_1, B_2 \subseteq X$ non-empty. But then Lemma 2.2.1 (3) implies that $B_1 \in \mathcal{F}$ or $B_2 \in \mathcal{F}$, in contradiction to the definition of A. So $A = \{a\}$ for some $a \in X$.

(2) ⇒ (3). Clearly, $\{Y \subseteq X \mid a \in Y\} \subseteq \mathcal{F}$ since $\mathcal{F}$ is closed under supersets, and $\mathcal{F} \subseteq \{Y \subseteq X \mid a \in Y\}$ since $\mathcal{F}$ does not contain the empty set.

(3) ⇒ (4): Clearly $\mathcal{F}$ contains a finite set; use Lemma 2.2.1 (2) to check that $\mathcal{F}$ is an ultrafilter.

(4) ⇒ (1): If $A \in \mathcal{F}$ is finite, then $B := \bigcap F$ is finite, and hence B is the intersection of finitely many elements in $\mathcal{F}$, and hence in $\mathcal{F}$ since $\mathcal{F}$ is a filter. This shows that $\mathcal{F}$ is principal. □

Are there non-principal ultrafilters?

LEMMA 2.2.3 (Ultrafilter Lemma). *Every filter $\mathcal{F}$ is contained in a ultrafilter.*

PROOF. Let $\mathcal{M}$ be the set of all filters on X that contain $\mathcal{F}$, partially ordered by containment. Note that unions of chains of filters in this partial order are again filters. By Zorn's lemma (Theorem A.2.1), $\mathcal{M}$ contains a maximal filter. □

Non-principal ultrafilters are also called *free ultrafilters*. In particular the Fréchet filter is contained in an ultrafilter, which must be free:

LEMMA 2.2.4. *An ultrafilter is free if and only if it contains the Fréchet filter.*

PROOF. Let $\mathcal{U}$ be a free ultrafilter on X and let $x \in X$. Either $\{x\} \in \mathcal{U}$ or $X \setminus \{x\} \in \mathcal{U}$. As $\mathcal{U}$ is free, $\{x\} \notin \mathcal{U}$ (Lemma 2.2.2). Hence, $X \setminus \{x\} \in \mathcal{U}$ for every $x \in X$. Let $F \subseteq X$ be finite. Then

$$X \setminus F = \bigcap_{x \in F} (X \setminus \{x\}) \in \mathcal{U}\,.$$

Now let $\mathcal{U}$ be a principal ultrafilter, i.e., there is $x \in X$ with $\{x\} \in \mathcal{U}$ (Lemma 2.2.2). Then the element $X \setminus \{x\}$ of the Fréchet filters is not in $\mathcal{U}$. □

Exercises.

(19) Show that a set of subsets of a set X can be extended to an ultrafilter if and only if it has the FIP.
(20) Show that a set $\mathcal{F}$ of subsets of a set X can be extended to a free ultrafilter if and only if the intersection of every finite subset of $\mathcal{F}$ is infinite.
(21) Show that every filter $\mathcal{F}$ on a set X is the intersection of all ultrafilters on X that extend $\mathcal{F}$.
(22) Show that if $\mathcal{U}$ is a free ultrafilter on X, $S \in \mathcal{U}$, and $T \subseteq X$ is such that the symmetric difference $S \Delta T$ is finite, then $S \in \mathcal{U}$.
(23) Show that there are $2^{2^{|X|}}$ many ultrafilters on an infinite set X. Hint: first show that there is a family $\mathcal{F}$ of $2^{|X|}$ subsets of X such that for any $A_1, \dots, A_n, B_1, \dots, B_n \in \mathcal{F}$

$$A_1 \cap \cdots \cap A_n \cap (X \setminus B_1) \cap \cdots \cap (X \setminus B_n) \neq \emptyset.$$

(24) True or false: if $\mathcal{U}$ and $\mathcal{V}$ are free ultrafilters on an infinite set X, is there is a permutation π of X such that $S \in \mathcal{U}$ if and only if $\pi(S) \in \mathcal{V}$?[1]

[1]Thanks to Lukas Juhrich for the idea for this exercise.

2.3. Ultraproducts

Let τ be a signature, let $\mathcal{U}$ be an ultrafilter on X, and for each $a \in X$ let $\underline{M}_a$ be a τ-structure. The idea of ultraproducts is to define an "average" structure of all the structures $\underline{M}_a$.

Definition 2.3.1. We write

$$\prod_{a \in X} \underline{M}_a / \mathcal{U}$$

for the τ-structure $\underline{M}$ whose domain are the equivalence classes of the equivalence relation $\sim$ defined on the set

$$\prod_{a \in X} M_a \coloneqq \{g \colon X \to \bigcup_{a \in X} M_a \mid \text{ for all } x \in X : g(a) \in M_a\}$$

as follows:

$$g \sim g' \quad :\Leftrightarrow \quad \{a \in X \mid g(a) = g'(a)\} \in \mathcal{U}.$$

The equivalence class of a function g with respect to the equivalence relation $\sim$ will be denoted by $[g]$. The relations and functions of $\underline{M}$ are defined as follows:

- for constant symbols $c \in \tau$:
$$c^{\underline{M}} \coloneqq [a \mapsto c^{\underline{M}_a}]$$
- for relation symbols $R \in \tau$ of arity k:
$$R^{\underline{M}}([g_1], \dots, [g_k]) :\Leftrightarrow \{a \in X \mid R^{\underline{M}_a}(g_1(a), \dots, g_k(a))\} \in \mathcal{U}$$
- for function symbols $f \in \tau$ of arity k:
$$f^{\underline{M}}([g_1], \dots, [g_k]) = [g_0] :\Leftrightarrow \{a \in X \mid f^{\underline{M}_a}(g_1(a), \dots, g_k(a)) = g_0(a)\} \in \mathcal{U}$$

It is straightforward (but tedious) to verify that this is indeed well defined, i.e., the interpretation of function and relation symbols is independent from the choice of the representatives[2].

Exercises.

(25) Show that Definition 2.3.1 is well-defined.

(26) Let $n \in \mathbb{N}$ and let $\underline{M}$ be a ultraproduct of finite structures each of which has at most n elements. Then $\underline{M}$ has at most n elements, too.

(27) (Exercise 1.31 in [**17**]). Let X be an index set and $\mathcal{U}$ an ultrafilter on X. Let $(\underline{A}_a)_{a \in X}$ and $(\underline{B}_a)_{a \in X}$ be families of τ-structures. If $\underline{A}_a$ can be embedded in $\underline{B}_a$ for all $a \in X$, show that $\prod_{a \in X} \underline{A}_a / \mathcal{U}$ can be embedded in $\prod_{a \in X} \underline{B}_a / \mathcal{U}$.

2.4. The Theorem of Łoś

Let $\{\underline{M}_a \mid a \in X\}$ be a family of τ-structures, and let $\mathcal{U}$ be an ultrafilter on X.

Theorem 2.4.1 (Łoś). *Let $\phi(\overline{x})$ be a first-order τ-formula and $\overline{[g]}$ be a tuple of elements of the ultraproduct $\underline{M} \coloneqq \prod_{a \in X} \underline{M}_a / \mathcal{U}$. Then*

$$\underline{M} \models \phi(\overline{[g]}) \quad \Leftrightarrow \quad \{a \in X \mid \underline{M}_a \models \phi(\overline{g(a)})\} \in \mathcal{U}$$

Proof. By induction over the syntactic form of first-order formulas.

- if ϕ is atomic and of the form $R(x_1, \dots, x_n)$, then the statement follows from the definition of $R^{\underline{M}}$. If ϕ contains terms then we have to do an additional induction over the recursive form of terms, which we omit.

[2]Note that the same definition works even if $\mathcal{U}$ is not an ultrafilter, but just a filter: in this case the resulting structure is called a *reduced product*.

- Suppose that the statement holds for ϕ and for ψ. Then

$$\begin{aligned}
& \underline{M} \models (\phi \wedge \psi)(\overline{[g]}) \\
\Leftrightarrow\ & \underline{M} \models \phi(\overline{[g]}) \text{ and } \underline{M} \models \psi(\overline{[g]}) && \text{(semantics of conjunction)} \\
\Leftrightarrow\ & \{a \in X \mid \underline{M}_a \models \phi(\overline{g(a)})\} \in \mathcal{U} \\
& \qquad \text{and } \{a \in X \mid \underline{M}_a \models \phi(\overline{g(a)})\} \in \mathcal{U} && \text{(inductive assumption)} \\
\Leftrightarrow\ & \{a \in X \mid \underline{M}_a \models \phi(\overline{g(a)})\} \cap \{a \in X \mid \underline{M}_a \models \psi(\overline{g(a)})\} \in \mathcal{U} && \text{($\mathcal{U}$ is filter)} \\
\Leftrightarrow\ & \{a \in X \mid \underline{M}_a \models (\phi \wedge \psi)(\overline{g(a)})\} \in \mathcal{U} && \text{(semantics of conjunction)}
\end{aligned}$$

- Suppose that the statement holds for ϕ. Then

$$\begin{aligned}
& \underline{M} \models \neg\phi(\overline{[g]}) \\
\Leftrightarrow\ & \underline{M} \not\models \phi(\overline{[g]}) && \text{(semantics of negation)} \\
\Leftrightarrow\ & \{a \in X \mid \underline{M}_a \models \phi(\overline{g(a)})\} \notin \mathcal{U} && \text{(inductive assumption)} \\
\Leftrightarrow\ & \{a \in X \mid \underline{M}_a \not\models \phi(\overline{g(a)})\} \in \mathcal{U} && \text{($\mathcal{U}$ is ultrafilter)} \\
\Leftrightarrow\ & \{a \in X \mid \underline{M}_a \models \neg\phi(\overline{g(a)})\} \in \mathcal{U} && \text{(semantics of negation)}
\end{aligned}$$

- Finally, suppose that the statement holds for the formula $\phi(x, \bar{y})$.

$$\begin{aligned}
& \underline{M} \models (\exists x.\phi)(\overline{[g]}) \\
& \Leftrightarrow \text{there is } [f] \in M \text{ with } \underline{M} \models \phi([f], \overline{[g]}) \\
& \Leftrightarrow \text{there is } [f] \in M \text{ with } \{a \in X \mid \underline{M}_a \models \phi(f(a), \overline{g(a)})\} \in \mathcal{U} \\
& \Leftrightarrow \{a \in X \mid \underline{M}_a \models (\exists x.\phi)(\overline{g(a)})\} \in \mathcal{U}
\end{aligned}$$

To see the final equivalence, note that for every $[f] \in M$

$$\{a \in X \mid \underline{M}_a \models \phi(f(a), \overline{g(a)})\} \subseteq \{a \in X \mid \underline{M}_a \models (\exists x.\phi)(\overline{g(a)})\}.$$

Conversely, there is $[f] \in M$ with $\supseteq$. Define $f\colon X \to \bigcup_a M_a$ as follows. For every $a \in X$ with $\underline{M}_a \models (\exists x.\phi)(\overline{g(a)})$ choose a witness $h_a \in M_a$ for x. Define

$$f(a) := \begin{cases} h_a & \text{if } \underline{M}_a \models (\exists x.\phi)(\overline{g(a)}) \\ \text{arbitrary} & \text{otherwise.} \end{cases}$$

Then $\{a \in X \mid \underline{M}_a \models \phi(f(a), \overline{g(a)})\} = \{a \in X \mid \underline{M}_a \models (\exists x.\phi)(\overline{g(a)})\}$. □

The following can be obtained from the theorem of Łoś by applying it to sentences instead of formulas, and to *ultrapowers* instead of ultraproducts. An *ultrapower* is an ultraproduct where all the factors $\underline{M}_a$ are equal to a single structure $\underline{A}$; in this case we write $\underline{A}^X$ instead of $\prod_{a \in X} \underline{M}_a$.

COROLLARY 2.4.2. *Let* $\underline{M} := \underline{A}^X/\mathcal{U}$. *Then* $\mathrm{Th}(\underline{M}) = \mathrm{Th}(\underline{A})$.

The statement of the corollary is trivial if $\mathcal{U}$ is a principal filter (why?). The statement is interesting if $\mathcal{U}$ is free, as we see in the following example.

EXAMPLE 18. Consider $\underline{A} := (\mathbb{N}; +, *, <, 0, 1)$ and let $\mathcal{U}$ be a free ultrafilter on ω. There is an element u of $\underline{M} := \underline{A}^\omega/\mathcal{U}$ such that for every $n \in \mathbb{N}$

$$\underline{M} \models \underbrace{1 + \cdots + 1}_{n \text{ times}} < u$$

for example $u = [(1, 2, 3, \dots)] \in M$. This distinguishes $\underline{M}$ from $\underline{A}$. (Why is there no contradiction to Corollary 2.4.2?) △

EXAMPLE 19. Let $\mathcal{U}$ be a free ultrafilter on ω, and consider $\underline{A} := (\mathbb{R}; 0, 1, +, *, \leq)$. Then

$$\mathrm{Th}(\underline{A}) = \mathrm{Th}(\underline{A}^\omega/\mathcal{U})$$

How big is $\underline{A}^\omega/\mathcal{U}$?

$$|\underline{A}^\omega/\mathcal{U}| \leq |\mathbb{R}^{\aleph_0}| = (2^{\aleph_0})^{\aleph_0} = 2^{\aleph_0}$$

Note that $\underline{A}^\omega/\mathcal{U}$ has *'infinitesimal'* elements, i.e., elements x such that for all $n \in \mathbb{N}$

$$\underline{A}^\omega/\mathcal{U} \models 0 < x \wedge \underbrace{(1 + \cdots + 1)}_{n \text{ times}} * x < 1$$

for example

$$x = [(1, 1/2, 1/3, 1/4, \dots)].$$

Abraham Robinson used this idea to develop *'nonstandard analysis'* which gives a formal interpretation to the reasoning with infinitely small positive entities à la Leibniz, Euler, and Cauchy. △

2.5. Proof of the Compactness Theorem

Let T be a theory.

PROOF OF THE COMPACTNESS THEOREM, THEOREM 2.0.1. Clearly, if a theory T is satisfiable then all finite subsets of the theory are satisfiable as well. For the converse, assume that every finite subset S of T has a model $\underline{M}_S$. Let X be the set of all finite subsets of T. For $\phi \in T$, let

$$X_\phi := \{S \in X \mid \underline{M}_S \models \phi\}.$$

Then the set $\{X_\phi \mid \phi \in T\}$ has the FIP: if $\phi_1, \dots, \phi_n \in T$, let $S := \{\phi_1, \dots, \phi_n\} \in X$ and note that $\underline{M}_S \models \phi_i$ for all i, and $S \in X_{\phi_i}$. Hence, $S \in X_{\phi_1} \cap \cdots \cap X_{\phi_n}$ shows that $X_{\phi_1} \cap \cdots \cap X_{\phi_n}$ is non-empty, proving the FIP.

By Lemma 2.1.2 and the ultrafilter lemma (Lemma 2.2.3) there is an ultrafilter $\mathcal{U}$ that contains $\{X_\phi \mid \phi \in T\}$. Then $\underline{M} := \prod_{S \in X} \underline{M}_S/\mathcal{U}$ is a model of T: for $\phi \in T$ we have $X_\phi \in \mathcal{U}$ and $\underline{M} \models \phi$ because of the theorem of Łoś (Theorem 2.4.1). □

2.6. Proving the Ultrafilter Lemma with Compactness

Let $\mathcal{F}$ be a filter on a set X. We want to show that $\mathcal{F}$ is contained in an ultrafilter $\mathcal{U}$. Let $\tau = \{P\} \cup \{c_S \mid S \subseteq X\}$ be a signatur with a unary relation symbol P and constant symbols c_S. Informally, our idea is to construct a theory T such that for every model $\underline{M}$ of T we have

$$\underline{M} \models P(c_S) \Leftrightarrow S \in \mathcal{U}.$$

And here is such a theory.

$$\begin{aligned} T \; := & \{P(c_A) \Rightarrow P(c_B) \mid A \subseteq B \subseteq X\} \\ & \cup \{(P(c_A) \wedge P(c_B)) \Rightarrow P(c_{A \cap B}) \mid A, B \subseteq X\} \\ & \cup \{P(c_A) \Leftrightarrow \neg P(c_{X \setminus A}) \mid A \subseteq X\} \\ & \cup \{P(c_A) \mid A \in \mathcal{F}\} \end{aligned}$$

Claim: Every finite $T' \subseteq T$ is satisfiable. In T' there are only finitely many constant symbols $c_{S_1}, \dots, c_{S_n}$. Then there is $x \in \bigcap_{P(c_{S_i}) \in T'} S_i$ because $\mathcal{F}$ has the FIP. Let $\underline{M}'$ be a τ-structure with domain $\mathcal{P}(X)$ and

- $c_S^{\underline{M}'} := S$
- $P^{\underline{M}'}(S)$ iff $x \in S$.

We have $\underline{M}' \models T'$. The compactness theorem asserts the existence of a model $\underline{M}$ of T. Then

$$\mathcal{U} := \{S \subseteq X \mid P^{\underline{M}}(S)\}$$

is an ultrafilter that extends $\mathcal{F}$. $\square$

CHAPTER 3

The Löwenheim-Skolem Theorem

The theorem of Löwenheim and Skolem implies that if T is a theory with infinite models, then T has models of arbitrary infinite cardinality κ, unless the signature is larger than κ. The proof of the theorem of Löwenheim and Skolem and has two parts, *upwards* and *downwards* (Theorem 3.2.1). Going upwards is essentially an application of the compactness theorem. For going downwards, we need to introduce *elementary chains*.

3.1. Chains

We will need another important operation to build structures, namely the formation of *limits of chains*. *Chains* of τ-structures are a fundamental concept from model theory. Let I be a linearly ordered index set. Let $(\underline{A}_i)_{i\in I}$ be a sequence of τ-structures. Then $(\underline{A}_i)_{i\in I}$ is called a *chain* if $\underline{A}_i$ is a substructure of $\underline{A}_j$ for all $i < j$.

Definition 3.1.1. The *union* of a chain $(\underline{A}_i)_{i\in I}$ is a τ-structure $\underline{B} = \lim_{i\in I} \underline{A}_i$ defined as follows.

- The domain of $\underline{B}$ is $B := \bigcup_{i\in I} A_i$.
- For each relation symbol $R \in \tau$ put $a \in R^{\underline{B}}$ if $a \in R^{\underline{A}_i}$ for some $i \in I$.
- For each function symbol $f \in \tau$ put $f^{\underline{B}}(a) := b$ if $f^{\underline{A}_i}(a) = b$ for some $i \in I$.

Example 20. For each $n \in \mathbb{N}$, let

$$A_n := \{-n, -n+1, \dots, 0, 1, 2, \dots\} \subseteq \mathbb{Z}$$

and let $\underline{A}_n := (A_n; \leq)$ be the substructure of $(\mathbb{Z}; \leq)$ induced by A_n. Then $\mathrm{Th}(\underline{A}_n) = \mathrm{Th}(\underline{A}_0)$ for all $n \in \mathbb{N}$, but $\mathrm{Th}(\lim_{i\in I} \underline{A}_i)$ is different (why?). $\triangle$

Let $\underline{A}$ and $\underline{B}$ be τ-structures. A function $f: A \to B$ *preserves* a first-order τ-formula $\phi(x_1, \dots, x_n)$ if and only if for all $a_1, \dots, a_n \in A$

$$\underline{A} \models \phi(a_1, \dots, a_n) \text{ implies } \underline{B} \models \phi(f(a_1), \dots, f(a_n)).$$

Remark 3.1.2. This definition is compatible with the notion of preservation as introduced in Section 1.1.5, in the following sense: if $\underline{A}'$ is the $\tau \cup \{R\}$-expansion of $\underline{A}$ by the relation defined by ϕ in $\underline{A}$ (i.e., $R^{\underline{A}'} = \phi^{\underline{A}}$), and $\underline{B}'$ is the $\tau \cup \{R\}$ expansion of $\underline{B}$ by the relation defined by ϕ in $\underline{B}$, then f preserves ϕ if and only if f is a homomorphism from $\underline{A}'$ to $\underline{B}'$.

Functions that preserve all first-order formulas are called *elementary* (they are necessarily embeddings). If $\underline{B}$ is an extension of $\underline{A}$ such that the identity map from $\underline{A}$ to $\underline{B}$ is an elementary embedding, we say that $\underline{B}$ is an *elementary extension* of $\underline{A}$, and that $\underline{A}$ is an *elementary substructure* of $\underline{B}$, and write $\underline{A} \prec \underline{B}$. Note that $\underline{A} \prec \underline{B}$ if and only if $\underline{A}$ is a substructure of $\underline{B}$ and $\mathrm{Th}(\underline{B}_A) = \mathrm{Th}(\underline{A}_A)$.

Example 21. The structure $\underline{A}_{n+1}$ from Example 20 is not an elementary extension of $\underline{A}_n$ (why?). $\triangle$

Note that $\prec$ is a transitive relation. A chain is called an *elementary chain* if $\underline{A}_i \prec \underline{A}_j$ for all $i < j$.

LEMMA 3.1.3 (Tarski-Vaught elementary chain theorem). *Let $(\underline{A}_i)_{i\in I}$ be an elementary chain of τ-structures. Then $\underline{A}_i \prec \underline{B} := \lim_{i\in I} \underline{A}_i$ for each $i \in I$.*

PROOF. We have to show that for each $i \in I$, every τ-formula $\phi(x_1, \dots, x_k)$, and every $\bar{a} \in (A_i)^k$ we have

$$\underline{B} \models \phi(\bar{a}) \text{ iff } \underline{A}_i \models \phi(\bar{a})$$

First consider the case that ϕ is atomic of the form $R(x_1, \dots, x_k)$ for $R \in \tau$. We then have

$$\begin{aligned} \underline{B} \models R(\bar{a}) &\text{ iff } \bar{a} \in R^{\underline{B}} \\ &\text{ iff } \bar{a} \in R^{\underline{A}_i} \\ &\text{ iff } \underline{A}_i \models R(\bar{a}) \end{aligned}$$

Claim 1. If $t(x_1, \dots, x_k)$ is a τ-term and $\bar{a} \in (A_i)^k$, then $t^{\underline{B}}(\bar{a}) = t^{\underline{A}_i}(\bar{a})$. This can be shown by a straightforward induction on the structure of terms. Consequently, the statement is true for atomic formulas of the form $t = s$, for τ-terms t and s.

Claim 2. If the statement is true for ϕ, then the statement is true for $\neg\phi$:

$$\begin{aligned} \underline{B} \models \neg\phi(\bar{a}) &\text{ iff not } \underline{B} \models \phi(\bar{a}) \\ &\text{ iff not } \underline{A}_k \models \phi(\bar{a}) && \text{(by IA)} \\ &\text{ iff } \underline{A}_k \models \neg\phi(\bar{a}) \end{aligned}$$

The verification for formulas of the form $\phi_1 \wedge \phi_2$ is similarly straightforward. We therefore go straight to the existential quantifier.

Claim 3. If the statement is true for ϕ then the statement is also true for $\exists x_1 .\, \phi$.

$$\begin{aligned} & \underline{B} \models \exists x_1 .\, \phi(x_1, \dots, x_k)(b_2, \dots, b_k) \\ \text{iff} \quad & \text{there is } b_1 \in B \text{ such that } \underline{B} \models \phi(b_1, b_2, \dots, b_k) \\ \text{iff} \quad & \text{there is } b_1 \in A_j \text{ such that } \underline{A}_j \models \phi(b_1, b_2, \dots, b_k) \text{ for some } j > i \\ \text{iff} \quad & \underline{A}_j \models \exists x_1 .\, \phi(b_2, \dots, b_k) \\ \text{iff} \quad & \underline{A}_i \models \exists x_1 .\, \phi(b_2, \dots, b_k) \text{ since } \underline{A}_i \prec \underline{A}_j. \end{aligned}$$

□

We say that a τ-formula $\phi(x_1, \dots, x_n)$ is *satisfiable* in a τ-structure $\underline{B}$ if there are elements $b_1, \dots, b_n$ such that $\underline{B} \models \phi(b_1, \dots, b_n)$; in this case we also say that ϕ *is satisfied by $b_1, \dots, b_n$ in $\underline{B}$*.

LEMMA 3.1.4 (Tarski's Test). *Let $\underline{B}$ be a τ-structure and $A \subseteq B$. Then A is the domain of an elementary substructure of $\underline{B}$ if and only if every $(\tau \cup A)$-formula $\phi(x)$ which is satisfiable in $\underline{B}_A$ can be satisfied by an element of A.*

PROOF. First suppose that $\underline{A} \prec \underline{B}$ and that $\underline{B}_A \models \phi(b)$ for some $b \in B$. Then $\underline{B}_A \models \exists x.\phi(x)$, hence $\underline{A}_A \models \exists x.\phi(x)$. So there exists $a \in A$ such that $\underline{A}_A \models \phi(a)$. Thus $\underline{B}_A \models \phi(a)$.

Conversely, suppose that the condition of Tarski's test is satisfied. To verify that A is the domain of a substructure $\underline{A}$ of $\underline{B}$, let $a_1, \dots, a_k \in A$ and $f \in \tau$. Then $x = f(a_1, \dots, a_k)$ is satisfiable in $\underline{B}_A$, and hence there is an $a \in A$ with $\underline{B}_A \models a = f(a_1, \dots, a_k)$. To show that $\underline{A}$ is an elementary substructure of $\underline{B}$ it suffices to show that $\underline{A}_A$ and $\underline{B}_A$ satisfy exactly the same first-order sentences. We prove this by induction over the length of first-order sentences. The statement is clear for atomic sentences. The induction steps for $\psi = \neg\phi$ and $\psi = (\phi_1 \wedge \phi_2)$ follow immediately from the inductive assumption for the shorter sentences ϕ and ϕ_1, ϕ_2. It remains to consider the case $\psi = \exists x.\phi(x)$. If ψ holds in $\underline{A}_A$, there exists $a \in A$ such that $\underline{A}_A \models \phi(a)$. The induction hypothesis yields $\underline{B}_A \models \phi(a)$, thus $\underline{B}_A \models \psi$. For the converse suppose that $\underline{B}_A \models \psi$. Then $\phi(x)$ is satisfiable in $\underline{B}_A$ and by assumption we find $a \in A$ such that $\underline{B}_A \models \phi(a)$. By induction $\underline{A}_A \models \phi(a)$, and hece $\underline{A}_A \models \psi$. $\square$

We use Tarski's Test to construct small elementary substructures.

COROLLARY 3.1.5. *Let S be a subset of a τ-structure $\underline{B}$. Then $\underline{B}$ has an elementary substructure $\underline{A}$ containing S and of cardinality at most* $\max(|S|, |\tau|, \aleph_0)$.

PROOF. We construct A as the union of an ascending sequence $S = S_0 \subseteq S_1 \subseteq \cdots$ of subsets of B. Suppose that S_i is already defined. Choose an element $a_\phi \in B$ for every $(\tau \cup S_i)$-formula $\phi(x)$ which is satisfiable in $\underline{B}_{S_i}$, and define S_{i+1} to be S_i together with these a_ϕ. Then Tarski's Test implies that A is the domain of an elementary substructure. It remains to prove the bound on $|A|$. There are at most $\max(|\tau|, \aleph_0)$ many τ-formulas (Corollary A.4.3). Let $\kappa = \max(|S|, |\tau|, \aleph_0)$. There are κ many $(\tau \cup S)$-formulas; therefore $|S_1| \leq \kappa$. Inductively it follows for every i that $|S_i| \leq \kappa$. Finally we have $|A| \leq \kappa \cdot \aleph_0 = \kappa$. $\square$

Exercises.

(28) Suppose that $\underline{A}$ is an elementary substructure of $\underline{B}$, and $\underline{B}$ is an elementary extension of $\underline{C}$, and $A \subseteq C$. Is $\underline{A}$ an elementary substructure of $\underline{C}$?[1]

(29) Let A, B sets with $A \subseteq B$ and suppose that A is infinite. Show that $\underline{A} := (A; =)$ is an elementary substructure of $\underline{B} := (B; =)$.
Solution. We have to show that $\mathrm{id}_A \colon A \to B$ is an elementary embedding of $\underline{A}$ into $\underline{B}$, i.e., preserves all first-order formulas: if $\phi(x_1, \dots, x_n)$ is a first-order formula and $a_1, \dots, a_n \in A$, then $\underline{A} \models \phi(a_1, \dots, a_n)$, if and only if $\underline{B} \models \psi(a_1, \dots, a_n)$. We show this by induction on the size of ϕ. If ϕ is atomic, it must be of the form $x = y$, and in this case the statement is clear. The inductive step is clear if ϕ is of the form $\neg\psi$ or of the form $\phi_1 \wedge \phi_2$. The interesting case is that ϕ is of the form $\exists y.\psi(y, x_1, \dots, x_n)$. If $\underline{A} \models \phi(a_1, \dots, a_n)$ then clearly $\underline{B} \models \phi(a_1, \dots, a_n)$. Conversely, suppose that $\underline{B} \models \phi(a_1, \dots, a_n)$. Then there exists $b \in B$ such that $\underline{B} \models \psi(b, a_1, \dots, a_n)$. Since A is infinite there exists $c \in A \setminus \{a_1, \dots, a_n\}$. Note that $(\underline{B}, a_1, \dots, a_n)$ has an automorphism that exchanges b to c and fixes all other elements. The automorphism shows that $\underline{A} \models \psi(c, a_1, \dots, a_n)$, and hence $\underline{A} \models \phi(a_1, \dots, a_n)$. This concludes the induction.

(30) Is $(\mathbb{Q}; <)$ an elementary substructure of $(\mathbb{R}; <)$?

(31) Is $(\mathbb{Z}; <)$ an elementary substructure of $(\mathbb{Q}; <)$?

(32) (Exercise 2.2.4 in [**29**]) Let $\underline{A} = (\mathbb{R}; 0, <, f^{\underline{A}})$ where f is a unary function symbol. Let $\underline{B}$ be an elementary extension of $\underline{A}$. Call $b \in B$ an *infinitesimal* if $-\frac{1}{n} < x < \frac{1}{2}$ for all $n \in \mathbb{N}$. Show that if $f^{\underline{A}}(0) = 0$, then $f^{\underline{A}}$ is continuous in 0 if and only if for every elementary extension $\underline{B}$ of $\underline{A}$ the map $f^{\underline{B}}$ maps infinitesimals to infinitesimals.

(33) A formula is called *existential* if it is built from a quantifier-free formula by existential quantification (universal quantification is forbidden). Show that the class of all models of an existential sentence is closed under extensions.

(34) A formula is called $\forall\exists$ if it is built from an existential formula by universal quantification. Show that the class of all models of a $\forall\exists$ sentence is closed under unions of chains.

[1] Thanks to Anna Tölle for the idea for this exercise.

(35) Let ϕ be a formula in prenex normal form (see Exercise 11) whose quantifier-free part is written in conjunctive normal form (see Exercise 11). Then ϕ is called a *Horn formula* if each conjunct has at most one disjunct which is *positive*, i.e., an atomic formula. In other words, all but one of the disjuncts of each conjunct must be of the form $\neg\psi$ for some atomic formula ψ. Prove that the class of all models of a Horn sentence is closed under taking products.

3.2. Proving Löwenheim-Skolem

THEOREM 3.2.1 (Löwenheim-Skolem). *Let $\underline{A}$ be an infinite τ-structure, $S \subseteq A$, and κ an infinite cardinal.*

- *(Downwards) If $\max(|S|, |\tau|) \leq \kappa \leq |A|$ then $\underline{A}$ has an elementary substructure of cardinality κ containing S.*
- *(Upwards) If $\max(|A|, |\tau|) \leq \kappa$ then $\underline{A}$ has an elementary extension of cardinality κ.*

PROOF. Downwards: choose a set $S' \subseteq A$ that contains S and has cardinality κ, and apply Corollary 3.1.5.

Upwards: we first construct an elementary extension $\underline{A}'$ of $\underline{A}$ of cardinality *at least* κ. Choose a set ρ of new constants of cardinality κ. As $\underline{A}$ is infinite, all finite subsets of the theory $T := \mathrm{Th}(\underline{A}_A) \cup \{\neg(c = d) \mid c, d \in \rho, c \neq d\}$ are satisfiable. By the compactness theorem (Theorem 2.0.1) T has a model $\underline{A}''$, which must have cardinality at least κ. Clearly, the τ-reduct $\underline{A}'$ of $\underline{A}''$ is an elementary extension of $\underline{A}$. Finally we apply downwards Löwenheim-Skolem to $\underline{A}'$ and $S = A$ and obtain an elementary substructure of $\underline{A}'$ of cardinality exactly κ, which is an elementary extension of $\underline{A}$ (see Exercise 28). $\square$

Note that in part (1), the assumption $\kappa \geq \max(|S|, |\tau|)$ is certainly necessary in general.

COROLLARY 3.2.2. *A theory which has an infinite model has a model in every cardinality $\kappa \geq \max(|\tau|, \aleph_0)$.*

3.3. Vaught's test

The theorem of Löwenheim-Skolem implies that no theory with an infinite model can describe this model up to isomorphism. The best we can hope for is a unique model for a given cardinality.

DEFINITION 3.3.1. A theory T is called κ-categorical if all models of T of cardinality κ are isomorphic.

Note that the theory of a finite structure does not have infinite models, and hence is κ-categorical for every infinite κ. In Example 15 we have seen that $\mathrm{Th}(\mathbb{Q}; <)$ is $\aleph_0$-categorical. A second important example of an $\aleph_0$-categorical theory is the theory of the *countable random graph* $(\mathbb{V}; E)$.

EXAMPLE 22. The *countable random graph* is the (simple and undirected) graph that has the following *extension property*: for all finite disjoint subsets U, U' of $\mathbb{V}$ there exists a vertex $v \in \mathbb{V} \setminus (U \cup U')$ such that v is adjacent to all vertices in U and to no vertex in U'. The existence of such a graph will be shown in the following chapter. We claim that this graph has an ω-categorical theory: note that the extension property of $(\mathbb{V}; E)$ given above is a first-order property; a back-and-forth argument similar to the one given Example 15 shows that every countably infinite graph with this property is isomorphic to $(\mathbb{V}; E)$. $\triangle$

THEOREM 3.3.2 (Vaught's Test). *A κ-categorical theory T is complete if the following conditions are satisfied:*

- *T is consistent;*
- *T has no finite model;*
- *$|\tau| \leq \kappa$.*

PROOF. We have to show that all models $\underline{A}, \underline{B}$ of T have the same theory. Since $\underline{A}$ and $\underline{B}$ must be infinite, by Corollary 3.2.2 the theory $\mathrm{Th}(\underline{A})$ and $\mathrm{Th}(\underline{B})$ must have models $\underline{A}'$ and $\underline{B}'$ of cardinality κ. By κ-categoricity of T, we have $\mathrm{Th}(\underline{A}') = \mathrm{Th}(\underline{B}')$, hence $\mathrm{Th}(\underline{A}) = \mathrm{Th}(\underline{A}') = \mathrm{Th}(\underline{B}') = \mathrm{Th}(\underline{B})$. $\square$

Exercises.

(36) Is $(\mathbb{Z}; <)$ $\aleph_0$-categorical?

CHAPTER 4

Fraïssé Amalgamation

A relational[1] structure $\underline{A}$ is called *homogeneous* (sometimes also called *ultra-homogeneous* [**19**]) if every isomorphism between finite substructures of $\underline{A}$ can be extended to an automorphism of $\underline{A}$. We have already encountered an example of a homogeneous structure.

Example 23. The back-and-forth argument we presented in Example 15 shows in particular that isomorphisms between finite substructures of $(\mathbb{Q}; <)$ can be extended to automorphisms of $(\mathbb{Q}; <)$. Thus, $(\mathbb{Q}; <)$ is homogeneous. △

A versatile tool to construct countable homogeneous structures from classes of finite structures is the *amalgamation technique* à la Fraïssé. We present it here for the special case of *relational structures*; this is all that is needed in the examples we are going to present. For a stronger version of Fraïssé-amalgamation for classes of structures that might involve function symbols, see [**19**].

4.1. The Age of a Structure

In the following, let τ be a countable relational signature. The *age* of a τ-structure $\underline{A}$ is the class of all finite τ-structures that embed into $\underline{A}$. A class $\mathcal{C}$ has the *joint embedding property (JEP)* if for any two structures $\underline{B}_1, \underline{B}_2 \in \mathcal{C}$ there exists a structure $\underline{C} \in \mathcal{C}$ that embeds both $\underline{B}_1$ and $\underline{B}_2$.

Proposition 4.1.1. *Let $\mathcal{C}$ be a class of τ-structures. Then $\mathcal{C}$ is the age of a (countable) relational structure if and only if $\mathcal{C}$ is closed under isomorphisms and substructures, has the JEP, and contains countably many isomorphism classes.*

Proof. Let $\underline{C}_0, \underline{C}_1, \ldots$ be a set of representatives for each isomorphism class in $\mathcal{C}$. Construct a chain $\underline{A}_0, \underline{A}_1, \ldots$ as follows. $\underline{A}_0 := \underline{C}_0$. If $\underline{A}_i$ has been chosen, find $\underline{B} \in \mathcal{C}$ such that $\underline{A}_i \hookrightarrow \underline{B}$ and $\underline{C}_{i+1} \hookrightarrow \underline{B}$. Let $\underline{A}_{i+1}$ be an extension of $\underline{A}_i$ isomorphic to $\underline{B}$. Then the age of $\underline{A} = \lim_{i \in \mathbb{N}} \underline{A}_i$ equals $\mathcal{C}$. □

4.2. Amalgamation Classes

A class $\mathcal{C}$ of τ-structures has the *amalgamation property (AP)* if for all $\underline{A}, \underline{B}_1, \underline{B}_2 \in \mathcal{C}$ and embeddings $e_1 \colon \underline{A} \hookrightarrow \underline{B}_1$ and $e_2 \colon \underline{A} \hookrightarrow \underline{B}_2$ there are a $\underline{C} \in \mathcal{C}$ and embeddings $f_i \colon \underline{B}_i \hookrightarrow \underline{C}$ for $i \in \{1, 2\}$ such that

$$f_1 \circ e_1 = f_2 \circ e_2.$$

To verify the amalgamation property of classes of τ-structures that are closed under isomorphism, a slightly different perspective is sometimes convenient. The *union* of two relational τ-structures $\underline{B}_1, \underline{B}_2$ is the τ-structure $\underline{C}$ with domain $B_1 \cup B_2$ and relations $R^{\underline{C}} := R^{\underline{B}_1} \cup R^{\underline{B}_2}$ for every $R \in \tau$. The intersection of $\underline{B}_1$ and $\underline{B}_2$ is defined analogously. Let $\underline{B}_1, \underline{B}_2$ be τ-structures such that $\underline{A}$ is a substructure of both

[1]The entire theory can be adapted to general signatures that might also contain function symbols; to keep the exposition simple, we restrict our focus to relational signatures in this section.

$\underline{B}_1$ and $\underline{B}_2$ and all common elements of $\underline{B}_1$ and $\underline{B}_2$ are elements of $\underline{A}$; note that in this case $\underline{A} = \underline{B}_1 \cap \underline{B}_2$. Then we call $\underline{B}_1 \cup \underline{B}_2$ the *free amalgam* of $\underline{B}_1, \underline{B}_2$ (over $\underline{A}$). More generally, a τ-structure $\underline{C}$ is an *amalgam of* $\underline{B}_1$ *and* $\underline{B}_2$ *over* $\underline{A}$ if for $i \in \{1,2\}$ there are embeddings f_i of $\underline{B}_i$ to $\underline{C}$ such that $f_1(a) = f_2(a)$ for all $a \in \underline{A}$. We refer to $(\underline{A}, \underline{B}_1, \underline{B}_2)$ as an *amalgamation diagram.*

Definition 4.2.1. An isomorphism-closed class $\mathcal{C}$ of τ-structures

- has the *free amalgamation property* if for any $\underline{B}_1, \underline{B}_2 \in \mathcal{C}$ the free amalgam of $\underline{B}_1$ and $\underline{B}_2$ is contained in $\mathcal{C}$;
- has the *amalgamation property* if every amalgamation diagram $(\underline{A}, \underline{B}_1, \underline{B}_2)$ has an amalgam $\underline{C} \in \mathcal{C}$;
- is an *amalgamation class* if $\mathcal{C}$ is countable up to isomorphism, has the amalgamation property, and is closed under taking substructures.

Note that since we only look at relational structures here (and since we allow structures to have an empty domain), the amalgamation property of $\mathcal{C}$ implies the joint embedding property.

Example 24. The class of all finite graphs has AP: an amalgam of $\underline{B}_1$ and $\underline{B}_2$ is $\underline{C} := \underline{B}_1 \cup \underline{B}_2$ (and $f_i \colon \underline{B}_i \hookrightarrow \underline{C}$ the identity). Similarly, for every relational signature τ, the class of all finite τ-structures has AP. △

A systematic source of amalgamation classes is the following proposition.

Proposition 4.2.2. *Let* $\underline{B}$ *be a homogeneous relational structure. Then* $\mathrm{Age}(\underline{B})$ *is an amalgamation class.*

Proof. Let $\underline{A}, \underline{B}_1, \underline{B}_2 \in \mathrm{Age}(\underline{B})$ so that for $i \in \{1,2\}$ there are embeddings $e_i \colon \underline{A} \hookrightarrow \underline{B}_i$. By definition, there are embeddings $g_i \colon \underline{B}_i \hookrightarrow \underline{B}$. Let $\underline{A}'$ be the substructure of $\underline{B}$ with domain $g_1(e_1(A))$. Then the restriction of $g_2 \circ e_2 \circ e_1^{-1} \circ g_1^{-1}$ to A is an embedding of $\underline{A}'$ into $\underline{B}$, and by the homogeneity of $\underline{B}$ can be extended to an automorphism a of $\underline{B}$. But then the substructure $\underline{C}$ of $\underline{B}$ with domain $a(g_1(B_1)) \cup g_2(B_2)$ is $\mathrm{Age}(\underline{B})$ and $a \circ g_1 \colon \underline{B}_1 \to \underline{B}$ and $g_2 \colon \underline{B}_2 \to \underline{B}$ are embeddings of $\underline{B}_1$ and $\underline{B}_2$ into $\underline{C}$ satisfying $a \circ g_1(e_1(a)) = g_2(e_2(a))$ for all $a \in A$, showing AP. □

Since we have already seen that $(\mathbb{Q}; <)$ is homogeneous, the age of $(\mathbb{Q}; <)$, which is the class of all finite linear orders, has AP; this can also bee seen directly.

4.3. Fraïssé's theorem

We prove a converse to Proposition 4.2.2.

Theorem 4.3.1 (Fraïssé [**12**,**13**]; see [**19**]). *Let* τ *be a countable relational signature and let* $\mathcal{C}$ *be an amalgamation class of* τ*-structures. Then there is a homogeneous and at most countable* τ*-structure* $\underline{C}$ *whose age equals* $\mathcal{C}$*. The structure* $\underline{C}$ *is unique up to isomorphism, and called the* Fraïssé-limit *of* $\mathcal{C}$*.*

Proof. We first prove uniqueness. This can be shown by a *back and forth argument*, generalising the argument in Example 15. Let $\underline{C}$ and $\underline{D}$ be countable homogeneous of the same age. We have to show that $\underline{C}$ and $\underline{D}$ are isomorphic, and construct the isomorphism $f \colon \underline{C} \to \underline{D}$ by a back-and-forth argument, similarly to the proof that $(\mathbb{Q}; <)$ is ω-categorical in Example 15. Let $C = \{c_1, c_2, \dots\}$ and $D = \{d_1, d_2, \dots\}$. Suppose f is already defined on a finite subset F of C.

- Going forth: Let $i \in \mathbb{N}$ be smallest so that $c_i \notin F$. Then there is $e \colon \underline{C}[F \cup \{c_i\}] \hookrightarrow \underline{D}$. Note that the finite structures $\underline{D}[e(F)]$ and $\underline{D}[f(F)]$ are isomorphic via $f \circ e^{-1}$. By the homogeneity of $\underline{D}$, this isomorphism can be extended to an automorphism a of $\underline{D}$. Then $a \circ e = f$ and we extend f by setting $f(c_i) := a(e(c_i))$. Clearly, the extension of f thus defined is an embedding of $\underline{C}[F \cup \{c_i\}]$ into $\underline{D}$ because it is the composition of the embedding e with the automorphism a.

- Going back: let $i \in \mathbb{N}$ be smallest so that $d_i \notin f[F]$. Analogously, find $c \in C$ so that the extension $f(c) := d_i$ is an isomorphism.

A structure $\underline{C}$ is called *weakly homogenous* if for all $\underline{B} \in \mathrm{Age}(\underline{C})$, substructure $\underline{A}$ of $\underline{B}$, and $e \colon \underline{A} \hookrightarrow \underline{C}$ there is $g \colon \underline{B} \hookrightarrow \underline{C}$ which extends e. Note that in the proof above, we only needed weak homogeneity of $\underline{C}$ and $\underline{D}$ to construct f.

We now prove the existence of the homogeneous structure $\underline{C}$ from the statement of the theorem. The structure $\underline{C}$ can be constructed as a union of a chain $(\underline{C}_i)_{i \in \mathbb{N}}$ of structures in $\underline{C}_i \in \mathcal{C}$ such that if $\underline{A}, \underline{B} \in \mathcal{C}$, with $\underline{A}$ substructure of $\underline{B}$ and $e \colon \underline{A} \hookrightarrow \underline{C}_i$ for some $i \in \mathbb{N}$, then there are $j \in \mathbb{N}$ and $g \colon \underline{B} \hookrightarrow \underline{C}_j$ which extends e. Note that $\underline{C}$ is weakly homogeneous, and hence homogeneous, by the comments above.

Also note that $\mathrm{Age}(\lim_i \underline{C}_i) = \mathcal{C}$. Here, the inclusion $\subseteq$ is clear. For the converse inclusion, first note that for every $\underline{A} \in \mathcal{C}$ there is $\underline{B} \in \mathcal{C}$ such that $\underline{A} \hookrightarrow \underline{B}$ and $\underline{C}_0 \hookrightarrow \underline{B}$ by the JEP. So $\underline{B} \hookrightarrow \underline{C}_j$ for some $j \in \mathbb{N}$, and hence $\underline{A} \hookrightarrow \underline{B} \hookrightarrow \lim_i \underline{C}_i$.

Let P be a countable set of representatives for all $(\underline{A}, \underline{B}) \in \mathcal{C}^2$ such that $\underline{A}$ is a substructure of $\underline{B}$. Let $\alpha \colon \mathbb{N}^2 \to \mathbb{N}$ be a bijection such that $\alpha(i,j) \geq i$ for all i, j. Suppose $\underline{C}_k$ already constructed. Let $(\underline{A}_{k,i}, \underline{B}_{k,i}, f_{k,i})_{i \in \mathbb{N}}$ be a list of all triples $(\underline{A}, \underline{B}, f)$ where

- $(\underline{A}, \underline{B}) \in P$ and
- $f \colon \underline{A} \hookrightarrow \underline{C}_k$.

Let i, j be such that $k = \alpha(i,j)$. Construct $\underline{C}_{k+1}$ as amalgam of $\underline{C}_k$ and $\underline{B}_{i,j}$ so that $f_{i,j}$ extends to $\underline{B}_{i,j} \hookrightarrow \underline{C}_{k+1}$. □

EXAMPLE 25. Let $\mathcal{C}$ be the class of all finite partially ordered sets. Amalgamation can be shown by computing the transitive closure: when $\underline{C}$ is the free amalgam of $\underline{B}_1$ and $\underline{B}_2$ over $\underline{A}$, then the transitive closure of $\underline{C}$ gives an amalgam in $\mathcal{C}$. The Fraïssé-limit of $\mathcal{C}$ is called the *homogeneous universal partial order*. △

EXAMPLE 26. Let $\mathcal{C}$ be the class of all finite undirected graphs. It is even easier than in the previous examples to verify that $\mathcal{C}$ is an amalgamation class, since here the free amalgam itself shows the amalgamation property. The Fraïssé-limit of $\mathcal{C}$ is also known as the *countable random graph*, or also the *Rado graph*, and will be denoted by $(\mathbb{V}; E)$. △

EXAMPLE 27. Let $\mathcal{C}$ be the class of all finite *triangle-free graphs*, that is, all finite undirected graphs that do not contain K_3 as a subgraph. Again, we have the free amalgamation property. The Fraïssé-limit is up to isomorphism uniquely described as the triangle-free graph $\underline{A}$ such that for any finite disjoint $S, T \subset A$ such that S *is stable* (i.e., induces a graph with no edges; such a vertex subset is sometimes also called an *independent set*) there exists $v \in A \setminus (S \cup T)$ which is connected to all points in S, but to no point in T. △

We now introduce a convenient notation to describe classes of finite τ-structures. When $\mathcal{N}$ is a class of τ-structures, we say that a structure $\underline{A}$ is $\mathcal{N}$*-free* if no $\underline{B} \in \mathcal{N}$ embeds into $\underline{A}$. The class of all finite $\mathcal{N}$-free structures we denote by $\mathrm{Forb}(\mathcal{N})$.

EXAMPLE 28. Henson **[16]** used Fraïssé limits to construct 2^ω many homogeneous directed graphs. A *tournament* is a directed graph without self-loops such that for all pairs x, y of distinct vertices exactly one of the pairs (x,y), (y,x) is an arc in the graph. Note that for all classes $\mathcal{N}$ of finite tournaments, $\mathrm{Forb}(\mathcal{N})$ is an amalgamation class, because if $\underline{A}_1$ and $\underline{A}_2$ are directed graphs in $\mathrm{Forb}(\mathcal{N})$ such that $\underline{A} = \underline{A}_1 \cap \underline{A}_2$ is a substructure of both $\underline{A}_1$ and $\underline{A}_2$, then the free amalgam $\underline{A}_1 \cup \underline{A}_2$ is also in $\mathrm{Forb}(\mathcal{N})$.

Henson specified an infinite set $\mathcal{T}$ of tournaments $\underline{T}_1, \underline{T}_2, \dots$ with the property that $\underline{T}_i$ does not embed into $\underline{T}_j$ if $i \neq j$. Draw the family on the board. Note that this property implies that for two distinct subsets $\mathcal{N}_1$ and $\mathcal{N}_2$ of $\mathcal{T}$ the two sets $\mathrm{Forb}(\mathcal{N}_1)$ and $\mathrm{Forb}(\mathcal{N}_2)$ are distinct as well. Since there are 2^ω many subsets of the infinite set $\mathcal{T}$, there are also that many distinct homogeneous directed graphs; they are often referred to as *Henson digraphs.* △

Exercises.

(37) Which of the following classes of finite structures are amalgamation classes?
   (a) The class of all finite directed graphs?
   (b) The class of all finite equivalence relations?
   (c) The class of all finite *left-linear orders*? A partial order $(P; \leq)$ is called *left-linear* if for every x the set $\{y \mid y \leq x\}$ is linearly ordered by $\leq$.
   (d) The class of all finite *3-uniform hypergraphs*? One way to formally define *3-uniform hypergraphs* is to view them as relational structures with a single ternary relation R which only contains triples with pairwise distinct elements and is *fully symmetric*, i.e., if $(a, b, c) \in R$ then $(\pi(a), \pi(b), \pi(c)) \in R$ for every permutation π of $\{a, b, c\}$.
   (e) The class of all finite 3-uniform hypergraphs that do not embed a *tetrahedron*: a tetrahedron is a 3-uniform hypergraph with vertices $\{1, 2, 3, 4\}$ whose edge relation contains *all* triples of pairwise distinct elements.

(38) The *complement* of a graph $(V; E)$ is the graph

$$\big(V; V^2 \setminus (E \cup \{(a, a) \mid a \in V\})\big).$$

Show that the complement of a homogeneous graph is homogeneous.

(39) Let $(\mathbb{V}; E)$ be the Rado graph (Example 26). Show that if we remove finitely many vertices from $(\mathbb{V}; E)$, the resulting graph is isomorphic to the Rado graph. Formally: show that for every finite $F \subseteq \mathbb{V}$, the subgraph of $(\mathbb{V}; E)$ induced by $\mathbb{V} \setminus F$ is isomorphic to $(\mathbb{V}; E)$.

(40) Let $(\mathbb{V}; E)$ and F be as in the previous exercise. Let $(\mathbb{V}; E')$ be the graph obtained from $(\mathbb{V}; E)$ by flipping edges and non-edges between F and $\mathbb{V} \setminus E$. Formally,

$$\begin{aligned} E' := E \setminus \{&(a, b) \in E \mid a \in F \text{ if and only if } b \notin F\} \\ \cup \{&(a, b) \notin E \mid a \in F \text{ if and only if } b \notin F\}. \end{aligned}$$

Show that $(\mathbb{V}; E')$ is isomorphic to $(\mathbb{V}; E)$.

(41) Show the age $\mathcal{C}$ of a structure has the amalgamation property if and only if it has the *1-point amalgamation property*, i.e., if for all $\underline{A}, \underline{B}_1, \underline{B}_2 \in \mathcal{C}$ and embeddings $e_1 \colon \underline{A} \hookrightarrow \underline{B}_1$ and $e_2 \colon \underline{A} \hookrightarrow \underline{B}_2$ such that $|B_1| = |B_2| = |A| + 1$ there are a $\underline{C} \in \mathcal{C}$ and embeddings $f_i \colon \underline{B}_i \hookrightarrow \underline{C}$ for $i \in \{1, 2\}$ such that $f_1 \circ e_1 = f_2 \circ e_2$.

4.4. Strong Amalgamation

A *strong amalgam* of $\underline{B}_1, \underline{B}_2$ over $\underline{A}$ is an amalgam of $\underline{B}_1, \underline{B}_2$ over $\underline{A}$ where $f_1(B_1) \cap f_2(B_2) = f_1(A) (= f_2(A))$ (as in Section 4.2). A class $\mathcal{C}$ has the *strong amalgamation property* if all amalgamation diagrams have a strong amalgam in $\mathcal{C}$.

Examples:

- The class of all finite graphs.

- The age of $(\mathbb{Q};<)$.

The following simple example shows an amalgamation class that does *not* have the strong amalgamation property.

EXAMPLE 29. Let P be a unary relation symbol, and let $\mathcal{C}$ be the class of all finite $\{P\}$-structures where P contains at most one element. Then $\mathcal{C}$ is an amalgamation class, and the Fraïssé-limit is a countably infinite structure where P contains only one element. But $\mathcal{C}$ does not have the strong amalgamation property. $\triangle$

For strong amalgamation classes there is a powerful construction to obtain new strong amalgamation classes from known ones.

DEFINITION 4.4.1. Let $\mathcal{C}_1$ and $\mathcal{C}_2$ be classes of finite structures with disjoint relational signatures τ_1 and τ_2, respectively. Then the *generic superposition of* $\mathcal{C}_1$ *and* $\mathcal{C}_2$, denoted by $\mathcal{C}_1 * \mathcal{C}_2$, is the class of $(\tau_1 \cup \tau_2)$-structures $\underline{A}$ such that the τ_i-reduct of $\underline{A}$ is in $\mathcal{C}_i$, for $i \in \{1,2\}$.

The following lemma has a straightforward proof by combining amalgamation in $\mathcal{C}_1$ with amalgamation in $\mathcal{C}_2$.

LEMMA 4.4.2. *If* $\mathcal{C}_1$ *and* $\mathcal{C}_2$ *are strong amalgamation classes, then* $\mathcal{C}_1 * \mathcal{C}_2$ *is also a strong amalgamation class.*

When $\underline{A}_1$ and $\underline{A}_2$ are homogeneous structures whose ages have strong amalgamation, then $\underline{A}_1 * \underline{A}_2$ denotes the (up to isomorphism unique) Fraïssé-limit of the generic superposition of the age of $\underline{A}_1$ and the age of $\underline{A}_2$.

PROPOSITION 4.4.3. *For* $i = 1$ *and* $i = 2$, *let* $\underline{A}_i$ *be a homogeneous* τ_i*-structure whose age has strong amalgamation. Then the* τ_i*-reduct of* $\underline{A}_1 * \underline{A}_2$ *is isomorphic to* $\underline{A}_i$.

PROOF. A back-and-forth argument. $\square$

EXAMPLE 30. For $i \in \{1,2\}$, let $\tau_i = \{<_i\}$, let $\mathcal{C}_i$ be the class of all finite τ_i-structures where $<_i$ denotes a linear order, and let $\underline{A}_i$ be the Fraïssé limit of $\mathcal{C}_i$. Then $\underline{A}_1 * \underline{A}_2$ is known as the *random permutation* (see e.g. **[8]**). $\triangle$

Exercises.

(42) Let $\underline{D}$ be the tournament obtained from the directed cycle $\vec{C}_3$ of length three by adding a new vertex u, and adding the edges (u,v) for every vertex v of $\vec{C}_3$. Let $\underline{D}'$ be the tournament obtained from $\underline{D}$ by flipping the orientation of each edge. Show that $\mathrm{Forb}(\{\underline{D},\underline{D}'\})$, the class of all finite tournaments that embeds neither $\underline{D}$ nor $\underline{D}'$, is an amalgamation class.

(43) Let P be a unary relation symbol. Let $\mathcal{D}$ be the class of all finite $\{P,<\}$-structures $\underline{A}$ such that $<^{\underline{A}}$ is a linear order.
  (a) Show that $\mathcal{D}$ is an amalgamation class.
  (b) Let $\underline{B}$ be the Fraïssé-limit of the class $\mathcal{D}$, and define $E \subseteq B^2$ by $(u,v) \in E$ if
    - $u < v$ and $(u \in P \Leftrightarrow v \in P)$, or
    - $u > v$ and not $(u \in P \Leftrightarrow v \in P)$.

    Show that $(B;E)$ is a tournament.
  (c) Show that the class $\mathrm{Age}(B;E)$ equals the class of tournaments that can be obtained from tournaments $\underline{T}$ in $\mathrm{Age}(\mathbb{Q};<)$ by performing the following operation: pick $u \in T$ and reverse all edges between u and other elements of T (we 'switch edges at u').
  (d) Show that $(B;E)$ is homogeneous.

(e) Show that $\mathrm{Age}(B;E)$ equals the class $\mathcal{C}$ from Exercise 42.
(f) Show that $(B;E)$ is isomorphic to the tournament whose vertices are a countable dense subset $S \subseteq \mathbb{R}^2$ of the unit circle without antipodal points, and where the edges are oriented in clockwise order, i.e., put $((u_1,u_2),(v_1,v_2)) \in E$ if and only if $u_1v_2 - u_2v_1 > 0$.

(44) Show that there are permutation groups G_1, G_2 on a countably infinite set such that both G_1 and G_2 are isomorphic (as permutation groups) to $(\mathbb{Q};<)$, but $G_1 \cap G_2 = \{\mathrm{id}\}$.
(45) Construct a permutation group G on a set X with precisely $n!$ orbits of n-element subsets.
(46) Show that the random graph can be partitioned into two subsets so that both parts are isomorphic to the random graph. Show that the same is not true for all partitions of the random graph into two infinite subsets.
(47) Show that for *every* partition of the vertices of the Rado graph into two subsets, one of the two subsets induces a subgraph of the Rado graph which is isomorphic to the Rado graph.
(48) A permutation group G on a set D is called *n-transitive* if the componentwise action of G on D^n has precisely one orbit. Construct for every $n \in \mathbb{N}$ a permutation group on D which is n-transitive, but not $(n+1)$-transitive.
(49) Give an example of a homogeneous structure with a transitive automorphism group whose age does not have strong amalgamation.
(50) Let $\underline{A}$ be a homogeneous structure with finite relational signature. Show that the following are equivalent.
(a) There exists a structure $\underline{B}$ with finite relational signature such that $\mathrm{Aut}(\underline{B}) = \mathrm{Aut}(\underline{A})$.
(b) There exists a relation $R \subseteq A^n$ such that $\mathrm{Aut}(\underline{A}) = \mathrm{Aut}(A,R)$.
(c) There exists a structure $\underline{A}$ with finite relational signature such that $G = \mathrm{Aut}(\underline{A})$, and all relations in $\underline{A}$ have pairwise distinct entries.
(d) There exists a relation $R \subseteq A^n$ such that $\mathrm{Aut}(\underline{B}) = \mathrm{Aut}(A,R)$ and R has pairwise distinct entries.

Which implications between these items require the homogeneity of $\underline{A}$?

4.5. Further Examples: C-relations

Let $\underline{T}$ be a tree with vertex set T and with a distinguished vertex r, the *root* of $\underline{T}$. For $u,v \in T$, we say that *u lies below v* if the path from r to u passes through v. We say that *u lies strictly below v* if u lies below v and $u \neq v$. The *youngest common ancestor (yca)* of two vertices $u,v \in T$ is the node w such that both u and v lies below w and w has maximal distance from r.

Let $\mathcal{T}$ be the class of all finite rooted binary trees $\underline{T}$. The *leaf structure* $\underline{C}$ of a tree $\underline{T} \in \mathcal{T}$ with leaves L is the relational structure $(L;|)$ where $|$ is a ternary relation symbol, and $ab|c$ holds in $\underline{C}$ iff $\mathrm{yca}(a,b)$ lies below $\mathrm{yca}(b,c)$ in $\underline{T}$ (recall that $\mathrm{yca}(a,b)$ denotes the youngest common ancestor of a and b in a rooted tree $\underline{T}$. We also call $\underline{T}$ the *underlying tree* of $\underline{C}$.

Observation 4.5.1. *Any $\underline{T} \in \mathcal{T}$ can be recovered uniquely from its leaf structure.*

Let $\mathcal{C}$ be the class of all leaf structures for trees from $\mathcal{T}$.

Proposition 4.5.2. *The class $\mathcal{C}$ is an amalgamation class.*

Proof. Closure under isomorphisms is by definition. Closure under substructures is easy to verify. For the amalgamation property, let $\underline{B}_1, \underline{B}_2 \in \mathcal{C}$ be such that $\underline{A} = \underline{B}_1 \cap \underline{B}_2$ is a substructure of both $\underline{B}_1$ and $\underline{B}_2$. We want to show that there is an amalgam of $\underline{B}_1$ and $\underline{B}_2$ over $\underline{A}$ in $\mathcal{C}$. We inductively assume that the statement has been shown for all triples $(\underline{A}, \underline{B}'_1, \underline{B}'_2)$ where $B'_1 \cup B'_2$ is a proper subset of $B_1 \cup B_2$.

Let $\underline{T}_1$ be the rooted binary tree underlying $\underline{B}_1$, and $\underline{T}_2$ the rooted binary tree underlying $\underline{B}_2$. Let $\underline{B}_1^1 \in \mathcal{C}$ be the substructure of $\underline{B}_1$ induced by the vertices below the left child of $\underline{T}_1$, and $\underline{B}_1^2 \in \mathcal{C}$ be the substructure of $\underline{B}_1$ induced by the vertices below the right child of $\underline{T}_1$. The structures $\underline{B}_2^1$ and $\underline{B}_2^2$ are defined analogously for $\underline{B}_2$ instead of $\underline{B}_1$.

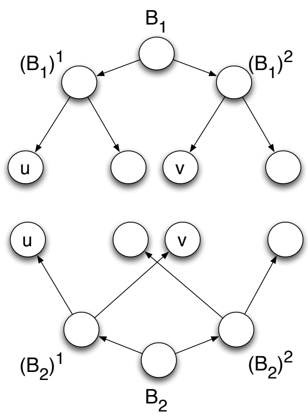

FIGURE 4.1. Illustration for the proof of Proposition 4.5.2.

First consider the case that there is a vertex u that lies in both $\underline{B}_1^1$ and $\underline{B}_2^1$, and a vertex v that lies in both $\underline{B}_2^1$ and $\underline{B}_1^2$ (see Figure 4.1 for an illustration). We claim that in this case no vertex w from $\underline{B}_2^2$ can lie inside $\underline{B}_1$: for otherwise, w is either in $\underline{B}_1^1$, in which case we have $uw|v$ in $\underline{B}_1$, or in $\underline{B}_1^2$, in which case we have $vw|u$ in $\underline{B}_1$. But since u, v, w are in A, this is in contradiction to the fact that $uv|w$ holds in $\underline{B}_2$. Let $\underline{C}' \in \mathcal{C}$ be the amalgam of $\underline{B}_1$ and $\underline{B}_2^1$ over $\underline{A}$, which exists by inductive assumption, and let $\underline{T}' \in \mathcal{T}$ be its underlying tree. Now let $\underline{T}$ be the tree with root r and $\underline{T}'$ as a left subtree, and the underlying tree of $\underline{B}_2^2$ as a right subtree. It is straightforward to verify that the leaf structure of $\underline{T}$ is in $\mathcal{C}$, and that it is an amalgam of $\underline{B}_1$ and $\underline{B}_2$ over $\underline{A}$ (via the identity embeddings).

Up to symmetry, the only remaining essentially different case we have to consider is that $B_1^1 \cup B_2^1$ and $B_1^2 \cup B_2^2$ are disjoint. In this case it is similarly straightforward to first amalgamate $\underline{B}_1^1$ with $\underline{B}_2^1$ and $\underline{B}_1^2$ with $\underline{B}_2^2$ to obtain the amalgam of $\underline{B}_1$ and $\underline{B}_2$; the details are left to the reader. □

Let $\underline{B}$ be the Fraïssé-limit of $\mathcal{C}$. The relation $|$ in $\underline{B}$ is closely related to so-called *C-relations*, following the terminology of [**1**]. C-relations became an important concept in model theory, see e.g. [**15**]. They are given axiomatically; the presentation here follows [**1**, **4**].

A ternary relation C is said to be a *C-relation* on a set L if for all $a, b, c, d \in L$ the following conditions hold:

C1 $C(a; b, c) \Rightarrow C(a; c, b)$;
C2 $C(a; b, c) \Rightarrow \neg C(b; a, c)$;
C3 $C(a; b, c) \Rightarrow C(a; d, c) \vee C(d; b, c)$;
C4 $a \neq b \Rightarrow C(a; b, b)$.

A C-relation is called *proper* if it satisfies

C5 $\forall b, c \exists a \colon C(a; b, c)$;
C6 $\forall a, b\big(a \neq b \Rightarrow \exists c(b \neq c \wedge C(a; b, c)\big)$.

A C-relation is called

C7 *dense* if it satisfies $C(a; b, c) \Rightarrow \exists e\,(C(e; b, c) \wedge C(a; b, e))$.
C8 *binary branching* if it satisfies

$$\forall x, y, z\,\big((x \neq y \vee x \neq z \vee y \neq z) \Rightarrow (C(x; y, z) \vee C(y; x, z) \vee C(z; x, y))\big).$$

PROPOSITION 4.5.3. *Every dense binary branching proper C-relation is isomorphic to the structure* $(B; \{(a, b, c) \in B^3 : a|bc \vee (a \neq b \wedge b = c)\})$ *(where* $\underline{B}$ *is the Fraïssé-limit of the class* $\mathcal{C}$ *as introduced above).*

PROOF. A back-and-forth argument. □

Exercises.

(51) Show that the class of all finite forests (i.e., undirected graphs without cycles) is *not* an amalgamation class.
(52) Prove Observation 4.5.1.
(53) Prove Proposition 4.5.3.

CHAPTER 5

Types

Loosely speaking, a *type* of a τ-structure $\underline{M}$ is a set of formulas that is satisfied by a real or by a 'virtual' element of $\underline{M}$, that is, an element of some structure that has the same theory as $\underline{M}$.

A (not necessarily finite) set Σ of formulas with free variables $x_1, \ldots, x_n$ is called *satisfiable over a structure* $\underline{A}$, or *realised in* $\underline{A}$, if there are elements $a_1, \ldots, a_n$ of $\underline{A}$ such that for all sentences $\phi \in \Sigma$ we have $\underline{A} \models \phi(a_1, \ldots, a_n)$. We say that Σ is *satisfiable* if there exists a structure $\underline{A}$ such that Σ is satisfiable over $\underline{A}$.

LEMMA 5.0.1. *A set Σ of formulas with free variables $x_1, \ldots, x_n$ is satisfiable if and only if all finite subsets of Σ are satisfiable.*

PROOF. Introduce new constant symbols $c_1, \ldots, c_n$. Then Σ is satisfiable if and only if $\Sigma(c_1, \ldots, c_n) := \{\phi(c_1, \ldots, c_n) \mid \phi(x_1, \ldots, x_n) \in \Sigma\}$ is satisfiable. Now apply the compactness theorem. □

For $n \geq 0$, an *n-type of a theory T* is a set p of formulas with free variables $x_1, \ldots, x_n$ such that $p \cup T$ is satisfiable. An *n-type of a structure* $\underline{A}$ is an n-type of the first-order theory of $\underline{A}$. Note that an n-type p of a theory T might or might not be realised in models $\underline{A}$ of T; if p is not realised in $\underline{A}$, then we also say that p is *omitted* in $\underline{A}$.

LEMMA 5.0.2. *Let $\underline{A}$ be a τ-structure and Σ a set of first-order τ-formulas with free variables $x_1, \ldots, x_n$. Then the following are equivalent.*

(1) Σ is an n-type of $\underline{A}$;
(2) every finite subset of Σ is realised in $\underline{A}$;
(3) $\underline{A}$ has an elementary extension that realises Σ.

PROOF. (3) $\Rightarrow$ (1): immediate.

(1) $\Rightarrow$ (2): Let Σ be an n-type of $\underline{A}$, i.e., there exists a model $\underline{B}$ of $\mathrm{Th}(\underline{A})$ and $\bar{b} \in B^n$ such that $\underline{B} \models \Sigma(\bar{b})$. Hence, if Ψ is a finite subset of Σ, then $\underline{B} \models \exists x_1, \ldots, x_n \bigwedge \Psi$ and since $\mathrm{Th}(\underline{A}) = \mathrm{Th}(\underline{B})$ we have $\underline{A} \models \exists x_1, \ldots, x_n \bigwedge \Psi$, so Ψ is realised in $\underline{A}$.

(2) $\Rightarrow$ (3): suppose that every finite subset Ψ of Σ is realised in $\underline{A}$. Then in particular every finite subset of $\Sigma \cup \mathrm{Th}(\underline{A}_A)$ is satisfiable, and hence $\Sigma \cup \mathrm{Th}(\underline{A}_A)$ has a model $\underline{B}$ by compactness. Then the τ-reduct of $\underline{B}$ is an elementary extension of $\underline{A}$ that realises Σ. □

An n-type p of a τ-theory T is *maximal*, or *complete*, if $T \cup p \cup \{\phi(x_1, \ldots, x_n)\}$ is unsatisfiable for any τ-formula $\phi \notin T \cup p$. In other words, for every τ-formula $\phi(x_1, \ldots, x_n)$, either $\phi \in p$ or $\neg\phi \in p$. We write $S_n(T)$ for the set of all complete n-types of T. The set of all first-order formulas with free variables $x_1, \ldots, x_n$ satisfied by an n-tuple $\bar{a} = (a_1, \ldots, a_n)$ in $\underline{A}$ is a maximal type of $\underline{A}$, and called *the type of $\bar{a}$ in $\underline{A}$*.

EXAMPLE 31. The structure $(\mathbb{Q}; <)$ has precisely one complete 1-type and precisely three complete 2-types. This follows easily from the homogeneity of $(\mathbb{Q}; <)$ (Exercise 55 and Example 23; also see Exercise 55). △

Exercises.

(54) Show that a structure $\underline{A}$ realises all 1-types if and only if $\underline{A}$ realises all n-types.

(55) Let $\underline{A}, \underline{B}$ be a τ-structures and f an isomorphism between $\underline{A}$ and $\underline{B}$. Then every tuple $\bar{a} \in A^n$ has the same type in $\underline{A}$ as $f(\bar{a})$ in $\underline{B}$.

(56) Prove or disprove:
- The set $\{x > n \mid n \in \mathbb{N}\}$ is a type of $(\mathbb{Z}, <)_{\mathbb{N}}$.
- The set $\{x > 0\} \cup \{x < n \mid n \in \{1, 2, 3, \dots\}\}$ is a type of $(\mathbb{Z}, <)_{\mathbb{N}}$.
- The set $\{x > n \mid n \in \mathbb{N}\}$ is realised $(\mathbb{Q}, <)_{\mathbb{N}}$.
- The set $\{x > 0\} \cup \{x < 1/n \mid n \in \{1, 2, 3, \dots\}\}$ is a type of $(\mathbb{Q}, <)_{\mathbb{Q}}$.

(57) Show that $(\mathbb{R}; 0, +)$ has exactly two complete 1-types and $\aleph_0$ many complete 2-types, and that $(\mathbb{R}; 0, +, <)$ has exactly three complete 1-types and $2^{\aleph_0}$ many complete 2-types.

Solution to Exercise 57.

(a) The elements 0 and 1 clearly have a different 1-type in $(\mathbb{R}; 0; +)$. We claim that every complete 1-type of $(\mathbb{R}; 0, +)$ that contains the formula $x_1 > 0$ is realised by 1. Let $\phi \in p$. By Lemma 5.0.2 the type $\{\phi(x_1), x_1 > 0\} \subset p$ is realised by some element a of $(\mathbb{R}; 0, +)$. Then $x \mapsto x/a$ is an automorphism of $(\mathbb{R}; 0; +)$ and takes a to 1, showing that ϕ also holds in 1.

(b) For any rational number $r/s \in \mathbb{Q}$, where $r \in \mathbb{Z}$ and $s \in \mathbb{N} \setminus \{0\}$, the pair (r, s) realises a maximal 2-type $p_{r,s}$. Note that $p_{r,s}$ contains the formula

$$\underbrace{x_1 + \cdots + x_1}_{s \text{ times}} = \underbrace{x_2 + \cdots + x_2}_{r \text{ times}}. \tag{1}$$

If $p_{u,v}$, for $u \in \mathbb{Z}$ and $v \in \mathbb{N} \setminus \{0\}$, also contains this formula, then $s \cdot u = r \cdot v$, so $r/s = u/v$. This implies that there are at least $|\mathbb{Q}| = \aleph_0$ many complete 2-types.

Then there is the complete 2-type of $(0, 0)$ and the complete 2-type of $(0, 1)$; they are clearly distinct and distinct from all the types of the form $p_{r,s}$ for some $r \in \mathbb{Z}$ and $s \in \mathbb{N} \setminus \{0\}$. I claim that there is exactly one more complete 2-type of $(\mathbb{R}; 0; +)$, namely the type of $(1, i)$ where i is any irrational number. To see this, view $\mathbb{R}$ as a $\mathbb{Q}$-vector space (of dimension $2^{\aleph_0}$); we may choose a basis B_1 that contains the element i. Let p be a complete 2-type $(\mathbb{R}; 0; +)$ which does not contain the formula $x_1 = 0$, does not contain the formula $x_2 = 0$, and does not contain (1) for any $r, s \in \mathbb{N}$. By Lemma 5.0.2, p is realised by an element a of an elementary extension $\underline{R}$ of $(\mathbb{R}; 0; +)$; by the theorem of Löwenheim-Skolem we may choose an elementary extension of cardinality $2^{\aleph_0}$. The elementary extension can be viewed as a $\mathbb{Q}$-vector space of dimension $2^{\aleph_0}$; using the axiom of choice, we may choose a basis B_2 that contains a. Then there exists a bijection between the basis elements B_1 and B_2 that extends to a vector space isomorphism between $(\mathbb{R}; 0; +)$ and $\underline{R}$ which shows that i satisfies p.

(c) The structure $(\mathbb{R}; 0, +, <)$ has the three complete 1-types realised by -1, 0, and 1, and the argument that there are all the complete 1-types can be shown similarly as in part (a) of this solution.

(d) There are clearly at most $2^{\aleph_0}$ many 2-types in $(\mathbb{R}; 0, +, <)$ because the signature is countable. For $a \in \mathbb{R}$, let p_a be the 2-type of the pair $(a, 1)$ in $(\mathbb{R}; 0, +, <)$. I claim that if $a, b \in \mathbb{R}$ are such that $a \neq b$, then $p_a \neq p_b$. It then follows that there are $|\mathbb{R}| = 2^{\aleph_0}$ many complete 2-types in $(\mathbb{R}; 0, +, <)$. Without loss of generality, suppose that $a < b$. Choose a rational number

r/s, for $r \in \mathbb{Z}$ and $s \in \mathbb{N} \setminus \{0\}$, such that $a < r/s < b$. Then

$$\underbrace{x_1 + \cdots + x_1}_{s \text{ times}} < \underbrace{x_2 + \cdots + x_2}_{r \text{ times}}$$

holds for $(a, 1)$, but not for $(b, 1)$.

5.1. Stone Spaces

There is a natural topology on the set $S_n(T)$ of complete n-types of a τ-theory T. For a τ-formula $\phi(x_1, \ldots, x_n)$, define

$$[\phi] := \{p \in S_n(T) \mid \phi \in p\}.$$

The *Stone topology* on $S_n(T)$ is the topology generated by taking the sets $[\phi]$ as basic open sets (also recall Remark 2.0.2). Note that for complete types p, exactly one of ϕ and $\neg\phi$ is in p; hence, $[\phi] = S_n(T) \setminus [\neg\phi]$ is also closed; we refer to sets that are both closed and open as *clopen*.

LEMMA 5.1.1. *The Stone topology on $S_n(T)$ is compact and totally disconnected.*

PROOF. To show compactness, let $C = \{[\phi_i] \mid i \in I\}$ be a cover of $S_n(T)$ by basic open sets. Thus, if $\underline{B}$ is a model of T and $a \in B^n$ we have that $\underline{B} \models \phi_i(a)$ for some $i \in I$. This shows that $T \cup \{\neg\phi_i \mid i \in I\}$ is unsatisfiable. So by compactness of first-order logic (Theorem 2.0.1) I must have a finite subset F such that $T \cup \{\neg\phi_i \mid i \in F\}$ is unsatisfiable. In other words, for every $p \in S_n(T)$ there is an $i \in F$ such that $\phi_i \in p$, which implies that $\{[\phi_i] \mid i \in F\}$ is a finite subcover of C.

To show total disconnectivity, let $p, q \in S_n(T)$ be distinct. Then there is a formula $\phi \in p$ such that $\neg\phi \in q$. Thus, $[\phi]$ is a basic clopen set separating p and q. $\square$

Exercises.

(58) Let $\underline{A}$ be a structure. Let T be the topology of A where the basic open sets are the subsets of A that can be defined in $\underline{A}$ by a first-order formulas $\phi(x)$ with parameters from A. Show that T is homeomorphic to $S_1(\mathrm{Th}(\underline{A}))$.

5.2. Saturated Structures

A structure $\underline{A}$ is *saturated* if, informally, 'as many types as possible' are realized in $\underline{A}$. Recall that $\underline{A}_B$, for $B \subseteq A$, denotes an expansion of $\underline{A}$ by constants for the elements of B. We refer to the n-types of $\underline{A}_B$ as the *n-types of $\underline{A}$ over B*. The set of all maximal n-types of $\underline{A}$ over B is denoted by $S_n^{\underline{A}}(B)$.

DEFINITION 5.2.1 (Saturation). For an infinite cardinal κ, a structure $\underline{A}$ is *κ-saturated* if $\underline{A}$ realises all 1-types over B for all $B \subseteq A$ with $|B| < \kappa$. We say that an infinite structure $\underline{A}$ is *saturated* if it is $|A|$-saturated.

EXAMPLE 32. The structure $(\mathbb{Q}; <)$ is saturated: if $q_1, \ldots, q_n \in \mathbb{Q}$, then the homogeneity of $(\mathbb{Q}; <)$ implies that $(\mathbb{Q}; <, q_1, \ldots, q_n)$ has precisely $2n + 1$ complete 1-types (see Example 15), and each of them is realised: we have n types realised by the elements $q_1, \ldots, q_n$, one type realised by an element smaller than all the q_i, one type realised by an element larger than all the q_i, and $n - 1$ types realised by elements that lie between some of the elements q_i. $\triangle$

The proof of the next result uses again a back and forth argument that we have seen already in Example 15.

THEOREM 5.2.2. *Let $\underline{A}$ and $\underline{B}$ be two saturated structures with the same first-order theory and the same cardinality. Then $\underline{A}$ and $\underline{B}$ are isomorphic.*

Proof. Let $(a_\alpha)_{\alpha<\kappa}$ be an enumeration of A and $(b_\alpha)_{\alpha<\kappa}$ an enumeration of B. We inductively construct a sequence $(c_\alpha)_{\alpha<\kappa}$ of elements of B and $(d_\alpha)_{\alpha<\kappa}$ of elements of A such that for all $\beta < \kappa$

$$\mathrm{Th}(\underline{A}; (a_\alpha)_{\alpha<\beta}, (d_\alpha)_{\alpha<\beta}) = \mathrm{Th}(\underline{B}; (c_\alpha)_{\alpha<\beta}, (b_\alpha)_{\alpha<\beta}). \tag{2}$$

The base case $\beta = 0$ holds by the assumptions of the theorem. Suppose that $(c_\alpha)_{\alpha<\beta}$ and $(d_\alpha)_{\alpha<\beta}$ have already been constructed. Let p be the 1-type of a_β in $(\underline{A}; (a_\alpha)_{\alpha<\beta}, (d_\alpha)_{\alpha<\beta})$. By (2) and since $\underline{B}$ is saturated there exists $c_\beta \in B$ that realises p. We will prove that

$$\mathrm{Th}(\underline{A}; (a_\alpha)_{\alpha\leq\beta}, (d_\alpha)_{\alpha<\beta}) = \mathrm{Th}(\underline{B}; (c_\alpha)_{\alpha\leq\beta}, (b_\alpha)_{\alpha<\beta}).$$

Let ϕ be a first-order sentence such that $(\underline{A}; (a_\alpha)_{\alpha\leq\beta}, (d_\alpha)_{\alpha<\beta}) \models \phi$. Let $\psi(x_1)$ be the first-order formula obtained from ϕ by replacing c_β by the new variable x_1. Note that ψ lies in p, and since c_β realises p we have $(\underline{B}; (c_\alpha)_{\alpha<\beta}, (b_\alpha)_{\alpha<\beta}) \models \psi(c_\beta)$. Hence, $(\underline{B}; (c_\alpha)_{\alpha\leq\beta}, (b_\alpha)_{\alpha<\beta}) \models \phi$. Similarly, we use saturation of $\underline{A}$ to find d_β such that

$$\mathrm{Th}(\underline{A}; (a_\alpha)_{\alpha\leq\beta}, (d_\alpha)_{\alpha\leq\beta}) = \mathrm{Th}(\underline{B}; (c_\alpha)_{\alpha\leq\beta}, (b_\alpha)_{\alpha\leq\beta}).$$

At the end of the day, the map $f\colon A \to B$ defined by $f(a_\alpha) := c_\alpha$ for all $\alpha < \kappa$ is a homomorphism from $\underline{A}$ to $\underline{B}$, and the map $b_\alpha \mapsto d_\alpha$ is a homomorphism from $\underline{B}$ to $\underline{A}$ which is the inverse of f. $\square$

We now prove general results about the existence of κ-saturated structures. We start with a lemma about realisation of types.

Lemma 5.2.3. *Every structure $\underline{A}$ has an elementary extension $\underline{B}$ that realises all 1-types over A.*

Proof. **First proof:** Let $(p_\alpha)_{\alpha<\lambda}$ be an enumeration of $S_1^{\underline{A}}(A)$ where λ is an ordinal[1]. We construct an elementary chain

$$\underline{A}_A =: \underline{A}_0 \prec \cdots \prec \underline{A}_\beta \prec \cdots (\beta \leq \lambda)$$

such that p_α is realised in $\underline{A}_{\alpha+1}$. Suppose that $(\underline{A}_{\alpha'})_{\alpha'<\beta}$ is already constructed.

- β is limit ordinal. Define $\underline{A}_\beta := \lim_{\alpha<\beta} \underline{A}_\alpha$. Then $(\underline{A}_{\alpha'})_{\alpha'\leq\beta}$ is an elementary chain, using Tarski's chain lemma (Lemma 3.1.3).
- $\beta = \alpha + 1$. Every finite subset Ψ of p_α is realised in $\underline{A}_A$, and therefore also in $\underline{A}_\alpha$. By Lemma 5.0.2, $\underline{A}_\alpha$ has an elementary extension $\underline{A}_{\alpha+1}$ that realises p_α.

Second proof: For each 1-type p of $\underline{A}$ over A, we introduce a new constant symbol c_p. Let T be the set of all atomic sentences of $\underline{A}_A$ together with all the formulas $\phi(c_p)$ where p is a 1-type of $\underline{A}$ over A and $\phi \in p$. We will show that finite subsets F of T are satisfiable. Let p be a 1-type of $\underline{A}$ over A. By definition, $\mathrm{Th}(\underline{A}_A) \cup p$ has a model $\underline{B}$, which satisfies

$$\psi := \exists x \bigwedge_{\phi(c_p)\in F} \phi(x).$$

Therefore, $\underline{A}_A \models \psi$, i.e., $\underline{A}_A \models \bigwedge_{\phi(c_p)\in F} \phi(a)$ for some $a \in A$. Expanding $\underline{A}_A$ by $c_p := a$ for all 1-types p of $\underline{A}$ over A, we obtain a model of F.

So by compactness, T has a model $\underline{B}$. Since T contains the atomic sentences that hold in $\underline{A}_A$, the structure $\underline{B}$ is an elementary extension of $\underline{A}$. Also, for each type p of $\underline{B}$ over A, the structure $\underline{B}$ contains an element c_p satisfying p, by choice of T. $\square$

[1]By Theorem A.3.4 there is a well-ordering of $S_1^{\underline{A}}(A)$, and by Proposition A.3.2 there is an ordinal λ and a sequence $(p_\alpha)_{\alpha<\lambda}$ that enumerates the elements of $S_1^{\underline{A}}(A)$.

THEOREM 5.2.4. *Let $\underline{A}$ be a structure and κ an infinite cardinal. Then $\underline{A}$ has a κ-saturated elementary extension.*

PROOF. Build a chain $(\underline{A}_\alpha)_{\alpha<\kappa^+}$ of structures inductively as follows:

- $\underline{A}_0 := \underline{A}$.
- $\underline{A}_{\alpha+1}$ is an elementary extension of $\underline{A}_\alpha$ realising all 1-types over A_α. Such a structure exists by Lemma 5.2.3.
- If β is a limit ordinal then $\underline{A}_\beta := \lim_{\alpha<\beta} \underline{A}_\alpha$.

It follows by induction on α that $\underline{A}_\alpha$ is an elementary extension of $\underline{A}_\beta$ for all $\beta < \alpha$, using Tarski's elementary chain lemma at the limit ordinals. So $(\underline{A}_\alpha)_{\alpha<\kappa^+}$ is an elementary chain of models.

By Tarski's elementary chain lemma again, $\underline{B} := \lim_{\alpha<\kappa^+} \underline{A}_\alpha$ is an elementary extension of $\underline{A}$. We show that $\underline{B}$ is even κ^+-saturated (which implies κ-saturation). Let S be a subset of B of size less than κ^+. Then $S \subseteq A_\alpha$ for some $\alpha \leq \kappa$: otherwise, S contains for every $\gamma < \kappa^+$ an element from $A_{\gamma+1} \setminus A_\gamma$, so $\mathrm{cf}(\kappa^+) < \kappa^+$, in contradiction to Proposition A.4.7. By construction, $\underline{A}_{\alpha+1}$ realises all 1-types over $\underline{A}_\alpha$. Consequently, $\underline{B}$ realises all 1-types over S. □

By paying attention to the sizes of the structures that we build, one can show the following.

THEOREM 5.2.5. *Let τ be a signature and let $\kappa \geq |\tau|$ an infinite cardinal. Then every τ-structure of size at most 2^κ has a κ^+-saturated elementary extension of size 2^κ.*

PROOF. The statement follows from the following observations:

- In the proof of Lemma 5.0.2, the elementary extension of $\underline{A}$ in item (3) can be chosen to have the same cardinality as $\underline{A}$, by applying the Löwenheim-Skolem theorem (Theorem 3.2.1).
- In the first proof of Lemma 5.2.3, the structure $\underline{B}$ can be build to have cardinality $2^{|A|}$ (there are at most $2^{|A|}$ many subsets of A).
- The union of a chain of length at most κ^+ of models of cardinality 2^κ has cardinality 2^κ (see Theorem A.4.2). □

So if we assume the Generalised Continuum Hypothesis, then there are saturated models of size κ^+ for all cardinals κ. We particularly point out the following special case.

COROLLARY 5.2.6. *Let τ be an at most countable signature, and T be a satisfiable τ-theory. Then T has a $\aleph_1$-saturated model of cardinality $2^{\aleph_0}$.*

Hence, assuming the continuum hypothesis, every satisfiable theory with a countable signature has a saturated model! Note that for general T, some set-theoretic assumption is necessary for the existence of saturated models: if T has $2^{\aleph_0}$ many 1-types (take for instance $(\mathbb{Q};<)$ expanded by constants for all elements) then any $\aleph_0$-saturated model has size $2^{\aleph_0}$. Hence, if $\aleph_1 < 2^{\aleph_0}$, then there is no saturated model of size $\aleph_1$. Saturated structures can also be constructed using ultraproducts.

THEOREM 5.2.7. *Let $(\underline{A}_i)_{i<\omega}$ be a sequence of structures with a countable signature τ and let $\mathcal{U}$ a non-principal ultrafilter on ω. Then $\underline{B} := \prod_{i<\omega} \underline{A}_i/\mathcal{U}$ is $\aleph_1$-saturated.*

PROOF. Let S be a subset of B of cardinality strictly smaller than $\aleph_1$, i.e., of cardinality $\aleph_0$. Let p be a 1-type over S. Note that $p = \{\phi_1, \phi_2, \dots\}$ is countable by our assumption that $|\tau| \leq \aleph_0$. For each $n \geq 1$, define

$$X_n := \big\{i \in \{n, n+1, \dots\} \mid \underline{A}_i \models \exists x(\phi_1(x) \wedge \dots \wedge \phi_n(x))\big\}.$$

Note that $\{n, n+1, \dots\} \in \mathcal{U}$ since $\mathcal{U}$ is non-principal. Also note that $\underline{B}$ realises $\phi_1(x) \wedge \cdots \wedge \phi_n(x)$ by Lemma 5.0.2, and hence $\{i \in \mathbb{N} \mid \underline{A}_i \models \exists x(\phi_1(x) \wedge \cdots \wedge \phi_n(x))\} \in \mathcal{U}$ by Theorem 2.4.1. Thus, $X_n \in \mathcal{U}$ for every $n \geq 1$. We have $X_{n+1} \subseteq X_n$ and $\bigcap_{n \geq 1} X_n = \emptyset$. So for every $i \in X_1$ there exists a maximal $n(i)$ such that $i \in X_{n(i)}$. We define an element $f \in \prod_{i \in \omega} \underline{A}_i$ as follows. For $i \in X_1$, let $f(i) \in A_i$ be such that $\underline{A}_i \models \phi_1(f(i)) \wedge \cdots \wedge \phi_{n(i)}(f(i))$; if $i \in \mathbb{N} \setminus X_1$, let $f(i)$ be any element of A_i (using the axiom of choice). We claim that $[f]_{\mathcal{U}}$ realises p in $\prod_{i \in \omega} \underline{A}_i / \mathcal{U}$. By Łoś's theorem (Theorem 2.4.1) we have to show that for every n

$$\{i \in \mathbb{N} \mid \underline{A}_i \models \phi_n(f(i))\} \in \mathcal{U}.$$

It suffices to show that $X_n \subseteq \{i \in \mathbb{N} \mid \underline{A}_i \models \phi_n(f(i))\}$ since $X_n \in \mathcal{U}$. Let $i \in X_n$. Then $i \in X_1$ and $i \in X_{n(i)}$. Thus, $\underline{A}_i \models \phi_{n(i)}(f_i)$. This shows the claim and concludes the proof. $\square$

5.3. Omitting Types

Let p be an n-type over T. A formula $\phi \in p$ *isolates* p if for every formula $\psi \in p$ we have that $T \models \forall \bar{x}(\phi(\bar{x}) \Rightarrow \psi(\bar{x}))$. A type p of a theory T is called *principal* if it is isolated by some formula. In the terminology of Section 5.1, the principal types are precisely the isolated points in $S_n(T)$: if ϕ isolates p, then $[\phi] = \{p\}$.

If T is complete, then every principal type p is realised in every model of T. To see this, let $\underline{B}$ be a model of T and suppose that the formula ϕ isolates the type p. Since p is a type, $T \cup \exists \bar{x}.\phi(\bar{x})$ is satisfiable, and since T is complete we we have that $\exists \bar{x}.\phi(\bar{x})$ is contained in T. Therefore, $\underline{B} \models \exists \bar{x}.\phi(\bar{x})$. If $a \in B$ is such that $\underline{B} \models \phi(\bar{a})$, then $\bar{a}$ realises p. The omitting types theorem can be viewed as a converse of this observation.

Theorem 5.3.1 (Omitting types theorem). *Let τ be a countable signature, let T be a satisfiable τ-theory, and let p be a non-principal n-type of T. Then T has a countable model that omits p.*

Proof. Let $\rho := \{c_0, c_1, \dots\}$ be countably many constant symbols that are not contained in τ. We will inductively construct a sequence $\theta_0, \theta_1, \dots$ of $(\tau \cup \rho)$-sentences such that $T^* := T \cup \{\theta_0, \theta_1, \dots\}$ has model $\underline{B}$ with an elementary substructure $\underline{A}$ that omits p. To do this, we alternate between steps that make sure that $\underline{A}$ omits p, and steps that make sure that we can apply Tarski's test (Lemma 3.1.4) to find the elementary substructure $\underline{A}$ of $\underline{B}$.

Let $\phi_0, \phi_1, \dots$ be an enumeration of all $(\tau \cup \rho)$-formulas $\phi(x)$. Let $d_0, d_1, \dots$ be an enumeration of ρ^n.

- Stage 0. Let θ_0 be $\top$.
- Stage $s = 2i + 1$ (for Tarki's test). Choose $c \in \rho$ which does not occur in $T \cup \{\theta_0, \dots, \theta_{2i}\}$ and define $\theta_{2i+1} := (\exists x.\phi_i(x) \Rightarrow \phi_i(c))$. Clearly, $T \cup \{\theta_0, \dots, \theta_{2i+1}\}$ is satisfiable.
- Stage $s = 2i+2$ (for omitting p). Let $d_i = (e_1, \dots, e_n)$. Let $\psi(x_1, \dots, x_n)$ be the τ-formula obtained from $\theta_1 \wedge \cdots \wedge \theta_{2i+1}$ by replacing each occurrence of e_i by x_i and then replacing every other symbol $c \in \rho \setminus \{e_1, \dots, e_n\}$ occurring in θ by the variable x_c and existentially quantifying over x_c. Because p is non-principal, there is a formula $\phi(x_1, \dots, x_n) \in p$ such that $T \cup \{\psi, \neg\phi\}$ is satisfiable, i.e., there is a model $\underline{B}$ of T and $b \in B^n$ such that $\underline{B} \models \psi(b) \wedge \neg\phi(b)$. Define $\theta_{2i+2} := \neg\phi(d_i)$. Note that $T \cup \{\theta_0, \dots, \theta_{2i+2}\}$ is satisfiable: a model $\underline{A}$ can be obtained as the expansion of $\underline{B}$ where $d_i^{\underline{A}} = b$.

The thus constructed theory T^* has a model $\underline{B}$ by compactness because $T \cup \{\theta_1, \dots, \theta_s\}$ is satisfiable for each $s \in \mathbb{N}$. Let $\psi(v)$ be a $(\tau \cup \rho)$-formula such that

$T^* \models \exists v.\psi(v)$. Then there is an $i \in \mathbb{N}$ such that $\psi = \phi_i$ and at stage $2i+1$ we have added the axiom $\exists v.\psi(v) \Rightarrow \psi(c)$ for some $c \in \rho$. Hence, by Tarski's test (Lemma 3.1.4), $\{c^{\underline{B}} \mid c \in \rho\}$ is the domain of an elementary substructure $\underline{A}$ of $\underline{B}$. Note that $\underline{A}$ is countable.

We claim that $\underline{A}$ omits p. Let $a \in A^n$. Then there is an $i \in \mathbb{N}$ such that $a = d_i^{\underline{B}}$. At stage $2i+2$ we ensure that $\underline{B} \models \neg\phi_i(d_i)$ for some $\phi_i \in p$. Thus, $\bar{a}$ does not realise p. This concludes the proof. $\square$

We mention that with the same proof idea one can also construct models of T that omit finitely many types at once. By cleverly designing countably many stages in the construction we may even omit countably may types at once.

THEOREM 5.3.2 (Countable omitting types theorem). *Let τ be a countable signature, let T be a satisfiable τ-theory, and let $p_1, p_2, \ldots$ be non-principal types of T. Then T has a countable model that omits all of the p_i's.*

We say that a structure $\underline{B}$ is *atomic* if for every $a \in B^n$, the complete type of a in $\underline{B}$ is principal.

THEOREM 5.3.3 (Theorem 6.2.2 in **[19]**). *Let T be a complete satisfiable theory with countably many n-types for every $n \in \mathbb{N}$. Then T has a countable atomic model.*

PROOF. There are only countably many non-principal complete types in T, so by the countable omitting types theorem (Theorem 5.3.1) there is a countable model $\underline{B}$ of T that omits all of them. $\square$

LEMMA 5.3.4. *Let $\underline{A}$ and $\underline{B}$ be atomic countable structures with the same theory. Then $\underline{A}$ and $\underline{B}$ are isomorphic.*

PROOF. Back and forth. $\square$

Exercises.

(59) Spell out the proof of Lemma 5.3.4.

(60) (Exercise 13.1.5 in **[25]**) Show that every countable model $\underline{M}$ of Peano Arithmetic has an elementary extension $\underline{N}$ such that for all $a \in M$ and $b \in N \setminus M$ we have $a < b$ (such extensions are called *end extensions*).

CHAPTER 6

Countably Categorical Structures

A structure $\underline{B}$ is called *ω-categorical* if its first-order theory is ω-categorical. There are many equivalent characterisations of ω-categoricity; the most important one is in terms of the automorphism group of $\underline{B}$. In the following, let $\mathscr{G}$ be a set of permutations of a set X. We say that $\mathscr{G}$ is a *permutation group* if $\mathscr{G}$ contains the identity id_X, and for $g, f \in \mathscr{G}$ the functions $x \mapsto g(f(x))$ and $x \mapsto g^{-1}(x)$ are also in $\mathscr{G}$. For $n \geq 1$ the *orbit* of $(t_1, \ldots, t_n) \in X^n$ under $\mathscr{G}$ is the set $\{(\alpha(t_1), \ldots, \alpha(t_n)) \mid \alpha \in \mathscr{G}\}$.

PROPOSITION 6.0.1. *Let $\underline{A}$ be an atomic countable structure. If $c, d \in A^n$ have the same type, then c and d lie in the same orbit of* $\mathrm{Aut}(\underline{A})$.

PROOF. We prove that if $c, d \in A^n$ have the same type, then for any $a \in A$ we can pick $b \in A$ such that then $n+1$-tuples (a, c) and (b, d) have the same type. Since $\underline{A}$ is atomic, the $n+1$-type p of (a, c) is principal; let $\psi(x_0, x_1, \ldots, x_n)$ be a formula that isolates it. Then $\exists x_0.\psi$ holds on c, and since types are preserved by automorphisms (Exercise 55) we have that $\exists x_0.\psi$ also holds on d. This gives us an element $b \in A$ such that (b, d) satisfies ψ; since ψ isolates p, it follows that (a, c) and (b, d) have the same type. Since $\underline{A}$ is countable, the statement now follows from a back and forth argument just as in the proof of Theorem 5.2.2. □

DEFINITION 6.0.2. A permutation group $\mathscr{G}$ over a countably infinite set X is *oligomorphic* if $\mathscr{G}$ has only finitely many orbits of n-tuples for each $n \geq 1$.

THEOREM 6.0.3 (Engeler, Ryll-Nardzewski, Svenonius). *For a countably infinite structure $\underline{B}$ with countable signature, the following are equivalent:*

(1) $\underline{B}$ is ω-categorical;
(2) every type of $\underline{B}$ is principal;
(3) every model of $\mathrm{Th}(\underline{B})$ *is atomic;*
(4) $\underline{B}$ has finitely many complete n-types, for all $n \geq 1$;
(5) for each $n \geq 1$, there are finitely many inequivalent formulas with free variables $x_1, \ldots, x_n$ over $\underline{B}$;
(6) every model of $\mathrm{Th}(\underline{B})$ *is ω-saturated;*
(7) every relation that is preserved by $\mathrm{Aut}(\underline{B})$ *is first-order definable in $\underline{B}$;*
(8) the automorphism group $\mathrm{Aut}(\underline{B})$ *of $\underline{B}$ is oligomorphic.*

PROOF. We show the following implications:

$$(1) \Rightarrow (2) \Rightarrow (3) \Rightarrow (1)$$
$$(2) \Rightarrow (4) \Rightarrow (5) \Rightarrow (6) \Rightarrow (1)$$
$$\big((2) \wedge (3) \wedge (5)\big) \Rightarrow (7) \Rightarrow (8) \Rightarrow (5)$$

$(1) \Rightarrow (2)$. Suppose $\underline{B}$ has a non-principal type p. By the omitting types theorem (Theorem 5.3.1) a countable model where p is omitted. But $\mathrm{Th}(\underline{B})$ also a countable model where p is realised by the theorem of Löwenheim-Skolem (Theorem 3.2.1), so $\underline{B}$ is not ω-categorical.

(2) $\Rightarrow$ (3). If all types of $\underline{B}$ are principal then $\underline{B}$ is atomic; the same applies to all models of $\mathrm{Th}(\underline{B})$.

(3) $\Rightarrow$ (1). All countable atomic models with the same theory are isomorphic by Lemma 6.0.1.

(2) $\Rightarrow$ (4). Suppose that all types are principal and let $n \geq 1$. Then there exists a sequence of formulas $(\phi_i)_{i \in I}$ such that every n-type is isolated by one of those formulas. Then $\mathrm{Th}(\underline{B}) \cup \{\neg\phi_i \mid i \in I\}$ is unsatisfiable and hence by the compactness theorem (Theorem 2.0.1) there exists a finite $F \subseteq I$ such that $\mathrm{Th}(\underline{B}) \cup \{\neg\phi_i \mid i \in F\}$ is unsatisfiable. That is, in every model of $\mathrm{Th}(\underline{B})$, every n-tuple satisfies ϕ_i for some $i \in F$, which shows that there are finitely many complete n-types in $\underline{B}$.

(4) $\Rightarrow$ (5). Every n-type is described by the complete n-types that contain the n-type, so if there are finitely many complete n-types, there are finitely many n-types in $\underline{B}$. And this provides a finite upper bound for the number of formulas with free variables $x_1, \dots, x_n$.

(5) $\Rightarrow$ (6). Let $\underline{A}$ be a model of $\mathrm{Th}(\underline{B})$, let $a \in A^n$, and let p be a complete 1-type of $(\underline{A}, a)$. If there is a finite number of inequivalent first-order formulas $\phi(x_1, x_2, \dots, x_{n+1})$, the conjunction over all formulas such that $\phi(x_1, \bar{a}) \in p$ isolates p. So p is realised in $(\underline{A}, a)$. This shows that $\underline{A}$ is ω-saturated.

(6) $\Rightarrow$ (1) follows from the fact that all countable ω-saturated structures with the same theory are isomorphic (Theorem 5.2.2).

(2) $\wedge$ (3) $\wedge$ (5) $\Rightarrow$ (7). Let R be an n-ary relation that is preserved by $\mathrm{Aut}(\underline{B})$. The relation R is a union of orbits of n-tuples of $\mathrm{Aut}(\underline{B})$. It suffices to show that orbits are first-order definable: by assumption (5), there are only finitely many inequivalent first-order formulas, we can then define R by forming a finite disjunction. Since $\underline{B}$ is atomic (3), if two n-tuples have the same type, then there is an automorphism that maps one to the other, by Lemma 6.0.1. So types define orbits of n-tuples. Since n-types of $\underline{B}$ are principal (2), it follows that the orbits of n-tuples are first-order definable in $\underline{B}$.

(7) $\Rightarrow$ (8). Suppose that $\mathrm{Aut}(\underline{B})$ are infinitely many orbits of n-tuples, for some n. Then the union of any subset of the set of all orbits of n-tuples is preserved by all automorphisms of $\underline{B}$; but there are only countably many first-order formulas over a countable language, so not all the invariant sets of n-tuples can be first-order definable in $\underline{B}$.

(8) $\Rightarrow$ (5) is immediate since automorphisms preserve first-order formulas. □

The second condition in Theorem 6.0.3 provides another possibility to verify that a structure is ω-categorical. We again illustrate this with the structure $(\mathbb{Q}; <)$. The orbit of an n-tuple $(t_1, \dots, t_n)$ from $\mathbb{Q}^n$ with respect to the automorphism group of $(\mathbb{Q}; <)$ is determined by the weak linear order induced by $(t_1, \dots, t_n)$ in $(\mathbb{Q}; <)$. We write *weak linear order*, and not *linear order*, because some of the elements $t_1, \dots, t_n$ might be equal (that is, a *weak linear order* is a total quasiorder). The number of weak linear orders on n elements is bounded by n^n, and hence the automorphism group of $(\mathbb{Q}; <)$ has a finite number of orbits of n-tuples, for all $n \geq 1$.

Two structures $\underline{A}$ and $\underline{B}$ with the same domain $A = B$ are called *interdefinable* if every relation and operation of $\underline{A}$ is first-order definable in $\underline{B}$, and vice versa.

Corollary 6.0.4. *Two ω-categorical structures $\underline{A}$ and $\underline{B}$ with the same domain $A = B$ have the same automorphism group if and only if they are interdefinable.*

Exercises.

(61) Is the structure $(\mathbb{Z}; \{(x, y) : |x - y| \leq 2\})$ ω-categorical?

(62) Show that if $\underline{A}$ is countable, $\underline{B}$ is ω-categorical, and $\mathrm{Age}(\underline{A}) \subseteq \mathrm{Age}(\underline{B})$, then $\underline{A}$ embeds into $\underline{B}$.
(63) Prove: there exists a countable bipartite graph that embeds all countable bipartite graphs.
(64) Show that if $\underline{B}$ is ω-categorical, then the expansion $(\underline{B}, c)$ of $\underline{B}$ by the constant $c \in B$ is ω-categorical as well.
(65) Which of the implications in Theorem 6.0.3 also work without the assumption that the signature is countable, and which implications fail?
(66) Which of the implications in Theorem 6.0.3 also work without the assumption that the structure $\underline{B}$ is countable, and which implications fail?

6.1. Algebraicity

The model-theoretic notion of algebraic closure generalises the notion of algebraic closure in fields (see Chapter B).

Definition 6.1.1. Let $\underline{A}$ be a structure and let $B \subseteq A$. We say that $a \in A$ is *algebraic over* B if there exists a $(\tau \cup \{B\})$-formula $\phi(x)$ such that $\{b \in A \mid \underline{A}_B \models \phi(b)\}$ is finite and contains a. The *algebraic closure of* B *in* $\underline{A}$ is the set of elements of $\underline{A}$ that are algebraic over B, and is denoted by $\mathrm{acl}_{\underline{A}}(B)$. We say that $\underline{A}$ *has algebraicity* if $\mathrm{acl}_{\underline{A}}(B) \neq B$ for some finite $B \subseteq A$.

Remark 6.1.2. Let $\underline{A}$ be a countable ω-categorical structure and let $B \subseteq A$ be finite. Then the algebraic closure of B in $\underline{A}$ is the set of elements of $\underline{A}$ that lie in finite orbits in $\mathrm{Aut}(\underline{A})_{(B)}$; this is a direct consequence of Theorem 6.0.3 (also see Exercise 64). Also note that for all finite subsets B the set $\mathrm{acl}_{\underline{A}}(B)$ is finite (see Exercise 69).

Theorem 6.1.3 (See (2.15) in **[7]**). *Let* $\underline{C}$ *be a homogeneous structure. Then the age of* $\underline{C}$ *has strong amalgamation if and only if* $\underline{C}$ *has no algebraicity.*

Proof. We first show that strong amalgamation of $\mathrm{Age}(\underline{C})$ implies no algebraicity of $\underline{C}$. Let $A \subseteq C$ be finite, and $u \in C \setminus A$. We want to show that the orbit of u in $\mathrm{Aut}(\underline{C})_{(A)}$ is infinite. Let $n \in \mathbb{N}$, $\underline{A} := \underline{C}[A]$, and $\underline{B} := \underline{C}[A \cup \{u\}]$. Then there exists a strong amalgam $\underline{B}' \in \mathrm{Age}(\underline{C})$ of $\underline{B}$ with $\underline{B}$ over $\underline{A}$. We iterate this, taking a strong amalgam of $\underline{B}$ with $\underline{B}'$ over $\underline{A}$, showing that, because of homogeneity, there are n distinct elements in $C \setminus A$ that lie in the same orbit as u in $\mathrm{Aut}(\underline{C})_{(A)}$. Since $n \in \mathbb{N}$ and $u \in C \setminus A$ were chosen arbitrarily, the group $\mathrm{Aut}(\underline{C})_{(A)}$ has no finite orbits outside A.

For the other direction, we rely on the following lemma of Peter Neumann.

Lemma 6.1.4. *Let* G *be a permutation group on* D *without finite orbits, and let* $A, B \subset D$ *be finite. Then there exists a* $g \in G$ *with* $g(A) \cap B = \emptyset$.

Proof. The proof here is from Cameron **[7]**, and is a nested induction. The outer induction is on $|A|$. We assume the result for any set A' with $|A'| < |A|$. The induction base $A = \emptyset$ is trivial. Suppose for contradiction that no $g \in G$ with the required property exists.

Claim. For any C with $|C| \leq |A|$, there are only finitely many translates $g(A)$ of A that contain C. The proof of the claim is by induction on $|A| - |C|$. When $|A| - |C| = 0$ then the only translate of A that contains C is C, and the statement holds. So suppose that $|C| < |A|$ and that the claim holds for all C' with $|C'| > |C|$. By the outer induction hypothesis, we may assume that $C \cap B = \emptyset$. By the inner induction hypothesis, for each of the finitely many points $b \in B$, only finitely many translates of A contain $C \cup \{b\}$. So only finitely many translates of A contain C and have non-empty intersection with B. Since we assumed that every translate of A has non-empty intersection with B, we have shown the claim.

For $C = \emptyset$, the claim implies that A has only finitely many translates, a contradiction to the assumption that G has no finite orbits. $\square$

We now continue with the reverse implication of Theorem 6.1.3. Let $\underline{A}$ be without algebraicity, and let $(\underline{A}_0, \underline{B}_1, \underline{B}_2)$ be an amalgamation diagram with $\underline{A}_0, \underline{B}_1, \underline{B}_2 \in \mathrm{Age}(\underline{A})$. By homogeneity of $\underline{A}$ we can furthermore assume that $\underline{A}_0, \underline{B}_1, \underline{B}_2$ are substructures of $\underline{A}$; that is, the structure induced by $B_1 \cup B_2$ in $\underline{A}$ is an amalgam, but possibly not a strong one. Since $\underline{A}$ has no algebraicity, $\mathrm{Aut}(\underline{A})_{(A_0)}$ has no finite orbits outside A_0. By the lemma, there exists $g \in \mathrm{Aut}(\underline{A})_{(A_0)}$ such that $(B_1 \setminus A_0) \cap (B_2 \setminus A_0) = \emptyset$. Then the structure induced by $g(B_1 \cup B_2)$ is a strong amalgam of $\underline{B}_1$ and $\underline{B}_2$ over $\underline{A}_0$. $\square$

Exercises.

(67) Find an example of a countable ω-categorical structure with only one 1-type and algebraicity.

(68) Show that if a $\underline{A}$ is structure and $B \subseteq A$ is finite such that $\mathrm{acl}_{\underline{A}}(B)$ is infinite, then $\underline{A}$ is not ω-categorical.

(69) Show that the converse of the implication in the previous exercise is false.

6.2. Quantifier Elimination

A τ-theory T *has quantifier elimination* if every τ-formula $\phi(x_1, \dots, x_n)$ is modulo T equivalent to a quantifier-free τ-formula $\psi(x_1, \dots, x_n)$, i.e.,

$$T \models \forall x_1, \dots, x_n (\phi \Leftrightarrow \psi).$$

If a theory T has quantifier elimination in a reasonable language then this can be very useful when working with T.

In this context, our assumption that we allow $\bot$ and $\top$ as a first-order formula (denoting the empty and full 0-ary relation, respectively) becomes relevant; Hodges **[18]** does not make this assumption, and therefore has to distinguish between *quantifier-elimination* and what he calls *quantifier-elimination for non-sentences.*

Remark 6.2.1. Note that one can extend any theory T to a theory T' with quantifier elimination by replacing τ by a larger signature that additionally contains for every first-order formula $\phi(x_1, \dots, x_n)$ a relation symbol R_ϕ, and by adding to T' all the axioms

$$\forall x_1, \dots, x_n \big(R_\phi(x_1, \dots, x_n) \Leftrightarrow \phi(x_1, \dots, x_n)\big).$$

This transformation preserves many model-theoretic property: if T is complete, or κ-categorical, then so is T'.

An interesting source of theories with quantifier-elimination is the following lemma.

Lemma 6.2.2. *An ω-categorical structure $\underline{B}$ has quantifier elimination if and only if it is homogeneous.*

Proof. Suppose first that $\underline{B}$ has quantifier-elimination. Let $\bar{a} = (a_1, \dots, a_k)$ and $\bar{b} = (b_1, \dots, b_k)$ be k-tuples of elements of $\underline{B}$ such that the mapping that sends a_i to b_i, for $1 \leq i \leq k$, is an isomorphism f between the structures induced by $\{a_1, \dots, a_k\}$ and by $\{b_1, \dots, b_k\}$. Since $\underline{B}$ has quantifier-elimination, these two tuples satisfy the same first-order sentences. By Theorem 6.0.3, the orbit of $(a_1, \dots, a_k)$ is first-order definable, and hence $(b_1, \dots, b_k)$ lies in the same orbit as $(a_1, \dots, a_k)$. It follows that f can be extended to an automorphism of $\underline{B}$.

Now suppose that $\underline{B}$ is homogeneous, and let $\phi(x_1,\dots,x_k)$ be a first-order formula. By the theorem of Ryll-Nardzewski (Theorem 6.0.3), there are finitely many orbits $O_1,\dots,O_m$ of k-tuples that satisfy ϕ. Clearly, it suffices to show that each of those orbits can be defined by a quantifier-free formula. Let $a \in B^k$ be such that $\underline{B} \models \phi(a)$. We claim that the set of quantifier-free formulas that hold on $(a_1,\dots,a_k)$ defines the orbit of a over $\underline{B}$. To see this, let $(b_1,\dots,b_k)$ be another k-tuple that satisfies the same quantifier-free formulas as $(a_1,\dots,a_k)$. Then the mapping that sends a_i to b_i is a partial isomorphism, and by homogeneity can be extended to an automorphism of $\underline{B}$. Since automorphisms preserve first-order formulas, $(b_1,\dots,b_k)$ also satisfies ϕ, which proves the claim. □

Exercises.

(70) Prove that the theory T' constructed in Remark 6.2.1 indeed has quantifier elimination.

(71) Which of the following ω-categorical structures has quantifier elimination:
- the structure $(\mathbb{N} \times \{0,1\}; E)$ where $E = \big\{((a,b),(a,c)) \mid b \neq c\big\}$.
- the structure $(\mathbb{N} \times \{0,1,2\}; E)$ where
$$E = \big\{((a,b),(a,c)) \mid (b,c) \in \{(0,1),(1,0),(1,2),(2,1)\}\big\}.$$
- the structure $(\mathbb{Q}_0^+; <)$ where $\mathbb{Q}_0^+ := \{x \in \mathbb{Q} \mid x \geq 0\}$.
- $(\mathbb{Q}; O)$ where $O := \{(x,y,z) \in \mathbb{Q}^3 \mid x < y \vee x < z\}$.
- $(\mathbb{Q}; B)$ where $B := \{(x,y,z) \in \mathbb{Q}^3 \mid x < y < z \vee z < y < x\}$.

6.3. First-Order Interpretations

First-order interpretations are a powerful tool to derive new structures from known structures.

Definition 6.3.1. Let $\underline{A}$ and $\underline{B}$ be structures with the relational signatures τ and σ and let $d \in \mathbb{N}$. A *(first-order) interpretation of dimension d* of $\underline{B}$ in $\underline{A}$ is a partial surjection $I\colon A^d \to B$ (also called the *coordinate map*) such that for every relation R, say of arity k, defined by an atomic σ-formula ϕ in $\underline{B}$, the dk-ary relation
$$I^{-1}(R) := \big\{(a_1^1,\dots,a_d^1,\dots,a_1^k,\dots,a_d^k) \mid (I(a_1^1,\dots,a_d^1),\dots,I(a_1^k,\dots,a_d^k)) \in R\big\}$$
has a first-order definition ϕ_I in $\underline{A}$.

Since $x = y$ is always allowed as an atomic formula, there must in particular exist a τ-formula $=_I$ such that $=_I (x_{1,1},\dots,x_{1,d},x_{2,1},\dots,x_{2,d})$ holds if and only if $\big((x_{1,1},\dots,x_{1,d}),(x_{2,1},\dots,x_{2,d})\big)$ lies in the kernel of I.

In order to specify a σ-structure $\underline{B}$ with a first-order interpretation in a given τ-structure $\underline{A}$ up to isomorphism, it suffices to specify the interpreting formulas for the atomic σ-formulas of $\underline{A}$; in particular, if the signature of $\underline{A}$ is relational and finite, then an interpretation has a finite presentation.

We say that $\underline{B}$ is *interpretable in $\underline{A}$ with finitely many parameters* if there are $c_1,\dots,c_n \in A$ such that $\underline{B}$ is interpretable in the expansion of $\underline{A}$ by the constants c_i for all $1 \leq i \leq n$.

Example 33. The field of rational numbers $(\mathbb{Q}; 0, 1, +, *)$ has a 2-dimensional interpretation I in $(\mathbb{Z}; 0, 1, +, *)$. The interpretation is now given as follows.
- The formula $=_I (x_1,x_2,y_1,y_2)$ is $x_2 \neq 0 \wedge y_2 \neq 0 \wedge x_1 y_2 = y_1 x_2$;
- The formula $0_I(x_1,x_2)$ is $x_2 \neq 0 \wedge x_1 = 0$, the formula $1_I(x_1,x_2)$ is $x_2 \neq 0 \wedge x_1 = x_2$;
- The formula $+_I(x_1,x_2,y_1,y_2,z_1,z_2)$ is
$$x_2, y_2, z_2 \neq 0 \wedge z_2 * (x_1 * y_2 + y_1 * x_2) = z_1 * x_2 * y_2;$$

- The formula $*_I(x_1, x_2, y_1, y_2, z_1, z_2)$ is
$$x_2, y_2, z_2 \neq 0 \wedge x_1 * y_1 * z_2 = z_1 * x_2 * y_2. \qquad \triangle$$

EXAMPLE 34. *Allen's Interval Algebra* is a structure studied in artificial intelligence for temporal reasoning; it is defined via a 2-dimensional first-order interpretation I in $(\mathbb{Q}; <)$. The domain formula $\top_I(x, y)$ is $x < y$. Hence, the elements of $\mathfrak{A}$ can indeed be viewed as non-empty closed bounded intervals $[x, y]$ over $\mathbb{Q}$. The template $\mathfrak{A}$ contains for each inequivalent $\{<\}$-formula ϕ with four variables a binary relation R such that (a_1, a_2, a_3, a_4) satisfies ϕ if and only if $\big((a_1, a_2), (a_3, a_4)\big) \in R$. In particular, $\mathfrak{A}$ has relations for equality of intervals, containment of intervals, and so forth. $\triangle$

EXAMPLE 35. Let $\underline{G} = (V; E)$ be an undirected graph (viewed as a symmetric digraph). Then the *line graph* $L_{\underline{G}}$ of $\underline{G}$ is the (undirected) graph with vertex set
$$V(L_{\underline{G}}) := \big\{\{u, v\} \mid (u, v) \in E\big\}$$
and the edge set
$$E(L_{\underline{G}}) := \big\{(\{u, v\}, \{v, w\}) \mid \{u, v\}, \{v, w\} \in E\big\}.$$
The line graph has the 2-dimensional first-order interpretation $I \colon E \to V(L_{\underline{G}})$ in $\underline{G}$ given by $I(x, y) := \{x, y\}$:

- $I^{-1}\big(\{(u, u) \mid u \in V(L_{\underline{G}})\}\big)$ has the first-order definition
$$\begin{aligned} E(x_1, x_2) \wedge E(y_1, y_2) \wedge \big((x_1 = y_1 \wedge x_2 = y_2) \\ \vee (x_1 = y_2 \wedge x_2 = y_1)\big). \end{aligned}$$
- $I^{-1}\big(E(L_{\underline{G}})\big)$ has the first-order definition
$$\begin{aligned} E(x_1, x_2) \wedge E(y_1, y_2) \wedge \big((x_1 = y_1 \wedge x_2 \neq y_2) \vee (x_1 = y_2 \wedge x_2 \neq y_1) \\ \vee (x_2 = y_1 \wedge x_1 \neq y_2) \vee (x_2 = y_2 \wedge x_1 \neq y_1)\big). \end{aligned} \qquad \triangle$$

Many ω-categorical structures can be derived from other ω-categorical structures via first-order interpretations.

LEMMA 6.3.2. *Let $\underline{A}$ be an ω-categorical structure. Then every structure $\underline{B}$ that is first-order interpretable in $\underline{A}$ with finitely many parameters is ω-categorical.*

PROOF. By the theorem of Ryll-Nardzewski (Theorem 6.0.3) it suffices to show that the number $o(n)$ of orbits of n-tuples under $\mathrm{Aut}(\underline{B})$ is finite, for every n. If $\underline{B}$ is the expansion of $\underline{A}$ by a constant c, then $o_{\underline{B}}(n) \leq o_{\underline{A}}(n + 1)$ since the map that sends the orbit of t to the orbit of (c, t) is an injection (see Exercise 64). If $\underline{B}$ has a d-dimensional interpretation in $\underline{A}$ then $o_{\underline{B}}(n) \leq o_{\underline{A}}(dn)$ and hence is finite, too. $\square$

Note that in particular all first-order reducts of an ω-categorical structure and all expansions of an ω-categorical structure by finitely many constants are again ω-categorical.

LEMMA 6.3.3. *Let $\underline{B}$ be a structure with at least two elements. Then every finite structure has a first-order interpretation in $\underline{B}$.*

PROOF. Let $\underline{A}$ be a τ-structure with domain $\{1, \ldots, n\}$. The statement is trivial if $n = 1$; so let us assume that $n > 1$ in the following. Our first-order interpretation I of $\underline{A}$ in $\underline{B}$ is n-dimensional. For $k \in \{1, \ldots, n-1\}$, define
$$\rho_k(x_1, \ldots, x_n) := \left(x_k \neq x_{k+1} \wedge \bigwedge_{i=1}^{k} x_1 = x_i\right)$$
$$\rho_n := (x_1 = \cdots = x_n).$$

The domain formula of our interpretation is *true*. Equality is interpreted by the formula

$$=_I (x_1,\dots,x_n,y_1,\dots,y_n) \coloneqq \bigvee_{k<n} \big(\rho_k(x_1,\dots,x_n) \wedge \rho_k(y_1,\dots,y_n)\big).$$

Note that the equivalence relation defined by $=_I$ on A^n has exactly n equivalence classes. If $R \in \tau$ is m-ary, then the formula $R(x_1,\dots,x_m)_I$ is a disjunction of conjunctions with the nm variables $x_{1,1},\dots,x_{m,n}$. For each tuple $(t_1,\dots,t_m)$ from $R^{\underline{A}}$ the disjunction contains the conjunct

$$\bigwedge_{i\leq m} \rho_{t_i}(x_{i,1},\dots,x_{i,n}). \qquad \square$$

Composing interpretations. First-order interpretations can be composed. In order to conveniently treat these compositions, we first describe how an interpretation of a σ-structure $\underline{B}$ gives rise to interpreting formulas ψ_I for arbitrary σ-formulas $\psi(x_1,\dots,x_n)$. If $\psi(y_1,\dots,y_n)$ is atomic, then $\psi_I(y_{1,1},\dots,y_{1,d},\dots,y_{n,1},\dots,y_{n,d})$ has already been defined. If ψ is of the form $\exists y.\psi'(y)$ by, then ψ_I is $\exists y_1,\dots,y_d.\,\psi'_I(y_1,\dots,y_d)$. Note that if ψ defines the relation R in $\underline{B}$, then ψ_I defines $I^{-1}(R)$ in $\underline{A}$. For all d-tuples $a_1,\dots,a_n \in I^{-1}(B)$

$$\underline{B} \models \psi\big(I(a_1),\dots,I(a_n)\big) \;\Leftrightarrow\; \underline{A} \models \psi_I(a_1,\dots,a_n)\ .$$

Definition 6.3.4. Let $\underline{C}$, $\underline{B}$, $\underline{A}$ be structures with the relational signatures ρ, σ, and τ. Suppose that

- $\underline{C}$ has a d-dimensional interpretation I in $\underline{B}$, and
- $\underline{B}$ has an e-dimensional interpretation J in $\underline{A}$.

Then $\underline{C}$ has a natural de-dimensional first-order interpretation $I \circ J$ in $\underline{A}$: the domain of $I \circ J$ is the set of all de-tuples t in A such that (t,t) satisfies the τ-formula $(=_I)_J$, and we define

$$I \circ J\big(a_{1,1},\dots,a_{1,e},\dots,a_{d,1},\dots,e_{d,e}\big) \coloneqq I\big(J(a_{1,1},\dots,a_{1,e}),\dots,J(a_{d,1},\dots,e_{d,e})\big).$$

Let ϕ be a τ-formula which defines a relation R over $\underline{A}$. Then the formula $(\phi_I)_J$ defines in $\underline{A}$ the preimage of R under $I \circ J$.

Let I_1 and I_2 be two interpretations of $\underline{B}$ in $\underline{A}$ of dimension d_1 and d_2, respectively. Then I_1 and I_2 are called *homotopic*[1] if the relation $\{(\bar{x},\bar{y}) \mid I_1(\bar{x}) = I_2(\bar{y})\}$ of arity $d_1 + d_2$ is first-order definable in $\underline{A}$. Note that id_C is an interpretation of $\underline{C}$ in $\underline{C}$, called the *identity interpretation* of $\underline{C}$ (in $\underline{C}$).

Definition 6.3.5. Two structures $\underline{A}$ and $\underline{B}$ with an interpretation I of $\underline{B}$ in $\underline{A}$ and an interpretation J of $\underline{A}$ in $\underline{B}$ are called *mutually interpretable*. If both $I \circ J$ and $J \circ I$ are homotopic to the identity interpretation (of $\underline{A}$ and of $\underline{B}$, respectively), then we say that $\underline{A}$ and $\underline{B}$ are *bi-interpretable (via I and J)*.

Example 36. The directed graph $\underline{C} \coloneqq (\mathbb{N}^2; M)$ where

$$M \coloneqq \big\{((u_1,u_2),(v_1,v_2)) \mid u_2 = v_1\big\}$$

and the structure $\underline{D} \coloneqq (\mathbb{N}; =)$ are bi-interpretable. The interpretation I of $\underline{C}$ in $\underline{D}$ is 2-dimensional, the domain formula is true, and $I(u_1,u_2) = (u_1,u_2)$. The interpretation J of $\underline{D}$ in $\underline{C}$ is 1-dimensional, the domain formula is true, and $J(x,y) = x$. Both interpretations are clearly .

Then $J(I(x,y)) = z$ is definable by the formula $x = z$, and hence $I \circ J$ is homotopic to the identity interpretation of $\underline{D}$. Moreover, $I(J(u),J(v)) = w$ is definable by

$$M(w,v) \wedge \exists p\,\big(M(p,u) \wedge M(p,w)\big),$$

[1]We follow the terminology from **[2]**.

so $J \circ I$ is also homotopic to the identity interpretation of $\underline{C}$. $\triangle$

EXAMPLE 37. Let $\mathbb{I}$ be the set of all non-empty closed bounded intervals over $\mathbb{Q}$ (also see Example 34). Let I be the 2-dimensional interpretation of $(\mathbb{I}; m)$ in $(\mathbb{Q}; <)$ with domain formula $x < y$, mapping $(x, y) \in \mathbb{Q}^2$ with $x < y$ to the interval $[x, y] \in \mathbb{I}$. The formula $(y_1 = y_2)_I$ is true, and the formula $(m(y_1, y_2))_I$ has variables $x_1^1, x_2^1, x_1^2, x_2^2$ and is given by $x_2^1 = x_1^2$.

Let J be the 1-dimensional interpretation with the domain formula true and $J([x, y]) := x$. The formula $(x < y)_I$ is the formula

$$\exists u, v \left(m(u, x) \wedge m(u, v) \wedge m(v, y) \right) .$$

We show that $J \circ I$ and $J \circ I$ are homotopic to the identity interpretation. The relation $\left\{(x_1, x_2, y) \mid J(I(x_1, x_2)) = y\right\}$ has the definition $x_1 = y$. To see that the relation $R := \left\{(u, v, w) \mid I(J(u), J(v)) = w\right\}$ has a definition in $(\mathbb{I}; m)$, first note that the relation

$$\left\{(u, v) \mid u = [u_1, u_2], v = [v_1, v_2], u_1 = v_1\right\}$$

has the definition $\phi_1(u, v) = \exists w \left(m(w, u) \wedge m(w, v) \right)$ in $(\mathbb{I}; m)$. Similarly, $\{(u, v) \mid u = [u_1, u_2], v = [v_1, v_2], u_2 = v_2\}$ has a definition $\phi_2(u, v)$. Then the formula $\phi_1(u, w) \wedge \phi_2(v, w)$ is equivalent to a primitive positive formula over $(\mathbb{I}; m)$, and defines R. $\triangle$

PROPOSITION 6.3.6. *If $\underline{A}$ and $\underline{B}$ are bi-interpretable, then* $\mathrm{Aut}(\underline{A})$ *and* $\mathrm{Aut}(\underline{B})$ *are isomorphic as abstract groups.*

PROOF. Let I be a d-dimensional interpretation of $\underline{B}$ in $\underline{A}$, and let J be an e-dimensional interpretation of $\underline{A}$ in $\underline{B}$ that witness that $\underline{A}$ and $\underline{B}$ are bi-interpretable. Let $\alpha \in \mathrm{Aut}(\underline{A})$; define $\mu(\beta) \colon B \to B$ as follows. For $b \in B$, pick $(a_1, \ldots, a_d) \in A^d$ such that $I(a_1, \ldots, a_d) = b$. Define $\mu(\beta)(b) := I(\alpha(a_1), \ldots, \alpha(a_d))$; note that this value is independent from the choice of $(a_1, \ldots, a_d)$, and that $\mu(\beta)$ defines an automorphism of $\underline{B}$. Conversely, every automorphism β of $\underline{B}$ induces an automorphism $\nu(\beta)$ of $\underline{A}$. The assumption that $I \circ J$ is homotopic to the identity interpretation of $\underline{A}$ implies that $\nu(\mu(\alpha)) = \alpha$, and the assumption that $J \circ I$ is homotopic to the identity interpretation of $\underline{B}$ implies that $\mu(\nu(\beta)) = \beta$, which shows that μ is a bijection with inverse ν. It is straighforward to verify that μ maps id_A to id_B and that it preserves composition. $\square$

EXAMPLE 38. The structures $\underline{C} := \left(\mathbb{N}^2; \{(x, y), (u, v) \mid x = u\}\right)$ and $\underline{D} := (\mathbb{N}; =)$ are mutually interpretable, but *not* bi-interpretable. There is a interpretation I_1 of $\underline{D}$ in $\underline{C}$, and a interpretation of $\underline{C}$ in $\underline{D}$ such that $I_2 \circ I_1$ is homotopic to the identity interpretation. However, the two structures are not bi-interpretable. To see this, observe that $\mathrm{Aut}(\underline{C})$ has the non-trivial normal subgroup $N := \mathrm{Sym}(\mathbb{N})^{\mathbb{N}}$ and that $\mathrm{Aut}(\underline{C})/N$ is isomorphic to $\mathrm{Aut}(\underline{D})$. However, $\mathrm{Aut}(\underline{D}) = \mathrm{Sym}(\mathbb{N})$ has exactly three proper non-trivial normal subgroups **[27]**, none of which is the automorphism group of a structure. Therefore, Proposition 6.3.6 implies that $\underline{C}$ and $\underline{D}$ are not bi-interpretable. $\triangle$

DEFINITION 6.3.7. A structure $\underline{B}$ has *essentially infinite signature* if every relational structure $\underline{C}$ that is interdefinable with $\underline{B}$ has an infinite signature.

We show that the property of having essentially infinite signature is preserved by bi-interpretability.

PROPOSITION 6.3.8. *Let $\underline{B}$ and $\underline{C}$ be structures that are bi-interpretable. Then $\underline{B}$ has essentially infinite signature if and only if $\underline{C}$ has essentially infinite signature.*

Proof. Let τ be the signature of $\underline{B}$. Suppose that the interpretation I_1 of $\underline{C}$ in $\underline{B}$ is d_1-dimensional, and that the interpretation I_2 of $\underline{B}$ in $\underline{C}$ is d_2-dimensional. Let $\theta(x, y_{1,1}, \dots, y_{d_1,d_2})$ be the τ-formula that shows that $I_2 \circ I_1$ is homotopic to the identity interpretation of $\underline{B}$. That is, θ defines in $\underline{B}$ the $(d_1 d_2 + 1)$-ary relation that contains a tuple $(a, b_{1,1}, \dots, b_{d_1,d_2})$ iff

$$a = h_2\big(h_1(b_{1,1}, \dots, b_{1,d_2}), \dots, h_1(b_{d_1,1}, \dots, b_{d_1,d_2})\big).$$

We have to show that if $\underline{C}$ has a finite signature, then $\underline{B}$ is interdefinable with a structure $\underline{B}'$ with a finite signature. Let $\sigma \subseteq \tau$ be the set of all relation symbols that appear in θ and in all the formulas of the interpretation of $\underline{C}$ in $\underline{B}$. Since the signature of $\underline{C}$ is finite, the cardinality of σ is finite as well. We will show that there is a definition of $\underline{B}$ in the σ-reduct $\underline{B}'$ of $\underline{B}$.

Let ϕ be an atomic τ-formula with k free variables $x_1, \dots, x_k$. Then the σ-formula

$$\begin{aligned}\exists y^1_{1,1}, \dots, y^k_{d_1,d_2}\ \Big(&\bigwedge_{i \leq k} \theta(x_i, y^i_{1,1}, \dots, y^i_{d_1,d_2}) \\ &\wedge\ \phi_{I_1 I_2}(y^1_{1,1}, \dots, y^k_{1,d_2}, y^1_{2,d_2}, \dots, y^k_{2,d_2}, \dots, y^k_{d_1,d_2})\Big)\end{aligned}$$

is equivalent to $\phi(x_1, \dots, x_k)$ over $\underline{B}'$. Indeed, by the surjectivity of h_2, for every element a_i of $\underline{B}$ there are elements $c^i_1, \dots, c^i_{d_2}$ of $\underline{C}$ such that $h_2(c^i_1, \dots, c^i_{d_2}) = a_i$, and by the surjectivity of h_1, for every element c^i_j of $\underline{C}$ there are elements $b^i_{1,j}, \dots, b^i_{d_1,j}$ of $\underline{B}$ such that $h_1(b^i_{1,j}, \dots, b^i_{d_1,j}) = c^i_j$. Then

$$\begin{aligned}\underline{B} \models R(a_1, \dots, a_k)\ &\Leftrightarrow \underline{C} \models \phi_{I_2}(c^1_1, \dots, c^1_{d_2}, \dots, c^k_1, \dots, c^k_{d_2}) \\ &\Leftrightarrow \underline{B}' \models \phi_{I_1 I_2}(b^1_{1,1}, \dots, b^k_{1,d_2}, b^1_{2,d_2}, \dots, b^k_{2,d_2}, \dots, b^k_{d_1,d_2}). \qquad \square\end{aligned}$$

Exercises.

(72) Show the claim from Example 38 that $(\mathbb{N}^2; \{(x,y),(u,v) \mid x = u\})$ and $(\mathbb{N}; =)$ are mutually interpretable.

(73) Show that the structure $(\mathbb{N}^2; \{((x,y),(y,z)) \mid x,y,z \in \mathbb{N}\})$ is bi-interpretable with $(\mathbb{N}; =)$.

(74) Show that $(\mathbb{N}^2; \{((x,y),(x,z)) \mid x,y,z \in \mathbb{N}\}, \{((x,y),(z,y)) \mid x,y,z \in \mathbb{N}\})$ and $(\mathbb{N}; =)$ are *not* bi-interpretable.

(75) (*) Show that two finite structures $\underline{A}$ and $\underline{B}$ have isomorphic automorphism groups (as abstract groups) if and only if $\underline{A}$ and $\underline{B}$ are bi-interpretable.

CHAPTER 7

Model Completeness and Quantifier-Elimination

Quantifier elimination in ω-categorical structures is particularly easy to check, as we have seen in Section 6.2. Proving quantifier elimination in more complicated mathematical structures such as $(\mathbb{C}; +, *, 1)$ or $(\mathbb{R}; +, 1, *, <)$ is more difficult. In this chapter we first introduce a concept which can be seen as half-way to quantifier elimination, namely model completeness. We will also learn some basic lemmas about quantifier elimination and then apply them to prove that the two interesting mathematical structures above have quantifier elimination.

7.1. Preservation Theorems

Preservation theorems in model theory establish links between definability in (a syntactically restricted fragment of) a given logic with certain 'semantic' closure properties.

Definition 7.1.1. When T is a first-order theory and $\phi(\bar{x})$ and $\psi(\bar{x})$ are formulas, we say that ϕ and ψ are *equivalent modulo T* if $T \models \forall \bar{x}(\phi(\bar{x}) \Leftrightarrow \psi(\bar{x}))$.

7.1.1. The Homomorphism Preservation Theorem. A formula is called *existential positive* if it is built from conjunction, disjunction, and existential quantification, only (negation and universal quantification are forbidden). Note that every existential positive τ-formula can be written as a disjunction of *primitive positive formulas*, i.e., formulas of the form

$$\exists x_1, \dots, x_n(\psi_1 \wedge \dots \wedge \psi_m)$$

where $\psi_1, \dots, \psi_m$ are atomic τ-formulas. Such formulas are also called *unions of disjunctive queries* in database theory and the most frequently encountered database queries in practice.

Such formulas also play an important role in geometry. For example, note that if $S_1, \dots, S_k$ are convex subsets of $\mathbb{R}^n$, and S_0 has a primitive positive definition in $(\mathbb{R}; S_1, \dots, S_k)$, then S_0 is convex as well. The homomorphism preservation theorem links existential positive definability modulo T with the fundamental concept of preservation under homomorphisms between models of T.

Theorem 7.1.2 (Homomorphism Preservation Theorem). *Let T be a first-order theory. A first-order formula ϕ is equivalent to an existential positive formula modulo T if and only if ϕ is preserved by all homomorphisms between models of T.*

Note that here the assumption that $\perp$ is always part of first-order logic is important: the first-order formula $\exists x. x \neq x$ is preserved by all homomorphisms between models of T, but without $\perp$ it might not be equivalent to an existential positive formula modulo T (for instance when T is the empty theory).

In the proof of Theorem 7.1.2 we need the following property of saturated structures.

Lemma 7.1.3. *Let $\underline{A}$ and $\underline{B}$ be τ-structures, where $\underline{B}$ is $|A|$-saturated. Suppose that every primitive positive sentence that is true in $\underline{A}$ is also true in $\underline{B}$. Then there is a homomorphism from $\underline{A}$ to $\underline{B}$.*

Proof. Let $(a_i)_{0 \leq i < |A|}$ well-order A. We claim that for every $\mu \leq |A|$ there is a sequence $(b_i)_{i<\mu}$ of elements of B such that every primitive positive $(\tau \cup \{c_i \mid i < \mu\})$-sentence true on $(\underline{A}, (a_i)_{i<\mu})$ is true on $(\underline{B}, (f(a_i))_{i<\mu})$. The proof is by transfinite induction on μ.

- The base case, $\mu = 0$, follows from the hypothesis of the lemma.
- The inductive step for limit ordinals μ follows from the observation that a sentence can only mention a finite collection of constants, whose indices must all be less than some $\nu < \mu$.
- For the inductive step for successor ordinals $\mu = \nu^+ < |A|$, set

$$\Sigma := \big\{\phi(x) \mid \phi \text{ is a ep-}(\tau \cup \{c_i \mid i < \nu\})\text{-formula such that } (\underline{A}, (a_i)_{i<\nu}) \models \phi(a_\mu)\big\} .$$

  By inductive assumption $(\underline{B}, (b_i)_{i<\nu}) \models \exists x.\phi(x)$ for every $\phi \in \Sigma$. By compactness, since Σ is closed under conjunction, we have that Σ is a ep-1-type of $(\underline{B}, (b_i)_{i<\nu})$. Then Σ is realised by some element $b_\gamma \in B$ because $\underline{B}$ is ep-$|A|$-saturated. By construction we maintain that all primitive positive $(\tau \cup \{c_i \mid i < \nu^+\})$-sentences true on $(\underline{A}, (a_i)_{i<\gamma^+})$ are true on $(\underline{B}, (b_i)_{i<\gamma^+})$.

The function that maps a_i to b_i for all $i < |A|$ is a homomorphism from $\underline{A}$ to $\underline{B}$. $\square$

Proof of Theorem 7.1.2. It is clear that homomorphisms preserve existential positive formulas. For the converse, let ϕ be first-order, with free variables $x_1, \ldots, x_n$, and preserved by homomorphisms between models of T. Let τ be the signature of T and ϕ, and let $\bar{c} = (c_1, \ldots, c_n)$ be a sequence of constant symbols that do not appear in τ. If $T \cup \{\phi(\bar{c})\}$ is unsatisfiable, the statement is clearly true, so assume otherwise. Let Ψ be the set of all existential positive $(\tau \cup \{c_1, \ldots, c_n\})$-sentences ψ such that $T \cup \{\phi(\bar{c})\} \models \psi$. Let $\underline{A}$ be a model of $T \cup \Psi$. Let U be the set of all primitive positive sentences θ such that $\underline{A} \models \neg\theta$.

We claim that $T \cup \{\neg\theta \mid \theta \in U\} \cup \phi(\bar{c})$ is satisfiable. For otherwise, by compactness, there would be a finite subset U' of U such that $T \cup \{\neg\theta \mid \theta \in U'\} \cup \phi(\bar{c})$ is unsatisfiable. But then $\psi := \bigvee_{\theta \in U'} \theta$ is an existential positive sentence such that $T \cup \{\phi(\bar{c})\} \models \psi$, and hence $\psi \in \Psi$. This is in contradiction to the assumption that $\underline{A} \models \neg\theta$ for all $\theta \in U$. We conclude that there exists a model $\underline{B}$ of $T \cup \{\neg\theta \mid \theta \in U\} \cup \phi(\bar{c})$.

By Theorem 5.2.5, $\underline{A}$ has an elementary extension $\underline{A}'$ which is $|B|$-saturated. Every primitive positive $(\tau \cup \{c_1, \ldots, c_n\})$-sentence θ that is true in $\underline{B}$ is also true in $\underline{A}'$: for if otherwise θ were false in $\underline{A}'$, then it were also false in $\underline{A}$, and hence $\theta \in U$ in contradiction to the assumption that $\underline{B} \models \{\neg\theta \mid \theta \in U\}$. Hence, by Lemma 7.1.3, there exists a homomorphism from $\underline{B}$ to $\underline{A}'$. Since $\underline{B} \models \phi(\bar{c})$, and ϕ is preserved by homomorphisms between models of T, we have $\underline{A}' \models \phi(\bar{a})$.

We conclude that $T \cup \Psi \cup \{\neg\phi(\bar{c})\}$ is unsatisfiable, and again by compactness there exists a finite subset Ψ' of Ψ such that $T \cup \Psi' \cup \{\neg\phi(\bar{c})\}$ is unsatisfiable. Then $\bigwedge \Psi'$ is an existential positive sentence; let ψ be the formula obtained from this sentence by replacing for all $i \leq n$ all occurrences of c_i by x_i. Then $T \models \forall \bar{x}(\psi(\bar{x}) \Leftrightarrow \phi(\bar{x}))$, which is what we wanted to show. $\square$

The classical theorem of Łos-Tarski for preservation under embeddings of models of a theory is a direct consequence of the homomorphism preservation theorem.

COROLLARY 7.1.4 (Łos-Tarski; see e.g. Corollary in 5.4.5 of [**19**]). *Let T be a first-order theory. A first-order formula ϕ is equivalent to an existential formula modulo T if and only if ϕ is preserved by all embeddings between models of T.*

PROOF. Add for each atomic formula ψ a new relation symbol N_ψ to the signature of T, and add the sentence $\forall \bar{x}(N_\psi(\bar{x}) \Leftrightarrow \neg\psi(\bar{x}))$; let T' be the resulting theory. Then every existential formula ϕ is equivalent to an existential positive formula in T', which can be obtained from ϕ by replacing negative literals $\neg\psi(\bar{x})$ in ϕ by $N_\psi(\bar{x})$. Similarly, homomorphisms between models of T' must be embeddings. Hence, the statement follows from Theorem 7.1.2. $\square$

7.1.2. The Theorem of Chang-Łos-Suszko. Another important preservation theorem concerns definability by so-called $\forall\exists$-formulas; theories consisting of $\forall\exists$-sentences are also called *inductive* for reasons that will soon become clear. Such theories will play an important role in Section 7.5.

DEFINITION 7.1.5. A formula ϕ is called $\forall\exists$ if it is of the form $\forall y_1, \dots, y_m.\ \psi$ where ψ is existential.

EXAMPLE 39. The property of the random graph which is described in Example 26 can be formulated as a set of $\forall\exists$-sentences. Note that this property described the random graph up to isomorphism. It follows that the theory of the random graph is equivalent to a $\forall\exists$-theory. $\triangle$

Our next goal will be a preservation theorem that characterises whether a formula is equivalent to a $\forall\exists$ formula modulo a theory T (Theorem 7.1.8 below). We first need the following simple but important and frequently used lemma.

LEMMA 7.1.6. *Let $\underline{A}$ and $\underline{B}$ be τ-structures and suppose that every existential sentence true in $\underline{A}$ is also true in $\underline{B}$. Then there exists an elementary extension $\underline{C}$ of $\underline{B}$ and an embedding $g\colon \underline{A} \to \underline{C}$.*

PROOF. Let T be the set of all quantifier-free sentences in $\mathrm{Th}(\underline{A}_A)$. It suffices to show that $T \cup \mathrm{Th}(\underline{B}_B)$ has a model $\underline{C}'$, since then

- the map $c^{\underline{A}_A} \mapsto c^{\underline{C}'}$ embeds $\underline{A}$ into the τ-reduct $\underline{C}$ of $\underline{C}'$, and
- $\underline{C}$ is an elementary extension of $\underline{B}$.

If $T \cup \mathrm{Th}(\underline{B}_B)$ has no model, then by the compactness theorem there is a τ-formula ϕ such that $\phi(\bar{c}) \in T$ and $\{\phi(\bar{c})\} \cup \mathrm{Th}(\underline{B}_B)$ has no model, and in particular $\underline{B} \models \neg\exists\bar{y}.\phi(\bar{y})$. Since $\exists\bar{y}.\phi(\bar{y})$ is existential, the assumptions imply that $\underline{A} \models \neg\exists\bar{y}.\phi(\bar{y})$. This contradicts that $\underline{A}_A \models \phi(\bar{c})$. $\square$

Similarly, one can show the following.

LEMMA 7.1.7. *Let T be a first-order theory, and let $\underline{A}$ be a model of the $\forall\exists$-consequences of T. Then $\underline{A}$ can be extended to a model $\underline{B}$ of T such that every existential formula that holds on a tuple $\bar{a}$ in $\underline{B}$ also holds on $\bar{a}$ in $\underline{A}$.*

PROOF. Let $\underline{A}'$ be an expansion of $\underline{A}$ by constants such that each element of $\underline{A}'$ is denoted by a constant symbol. It suffices to prove that $T \cup \mathrm{Th}(\underline{A}')_{\mathrm{qf}} \cup \mathrm{Th}(\underline{A}')_\forall$ has a model $\underline{B}$. Suppose for contradiction that it were inconsistent; then by compactness, there exists a finite subset U of $\mathrm{Th}(\underline{A}')_\forall \cup \mathrm{Th}(\underline{A}')_{\mathrm{qf}}$ such that $T \cup U$ is inconsistent. Let ϕ be the conjunction over U where all new constant symbols are existentially quantified. Then $T \cup \{\phi\}$ is inconsistent as well. But $\neg\phi$ is equivalent to a $\forall\exists$ formula, and a consequence of T. Hence, $\underline{A} \models \neg\phi$, a contradiction. $\square$

THEOREM 7.1.8 (Chang-Łoś-Suszko Theorem). *A theory T is equivalent to a $\forall\exists$-theory if and only if the class of models of T is closed under forming unions of chains.*

Proof. Unions of chains clearly preserve $\forall\exists$-sentences. Conversely, suppose that the models of T are preserved by forming unions of chains. Let S be the set of all $\forall\exists^+$-sentences that are consequences of T. We show that S implies T. It suffices to show that every model of S is elementary equivalent to a chain $(\underline{B}_i)_{i<\omega}$ of models of T. To construct this chain, we define an elementary chain of models $(\underline{A}_i)_{i<\omega}$ of T such that there are

- embeddings $f_i \colon \underline{A}_i \to \underline{B}_i$, with $\underline{B}_i \models T$, such that for every tuple $\bar{a}_i$ of elements from $\underline{A}_i$ and every existential formula θ, if $\underline{B}_i \models \theta(f_i(\bar{a}_i))$, then $\underline{A}_i \models \theta(\bar{a}_i)$, and
- embeddings $g_i \colon \underline{B}_i \to \underline{A}_{i+1}$, such that $g_i \circ f_i$ is the identity on $\underline{A}_i$.

Let $\underline{A}_0$ be a countable model of T. To construct the rest of the sequence, suppose that $\underline{A}_i$ has been chosen. Since $\underline{A}_0$ is an elementary substructure of $\underline{A}_i$, in particular all the $\forall\exists$-consequence of S hold in $\underline{A}_i$. By Lemma 7.1.7, the structure $\underline{A}_i$ can be extended to a model $\underline{B}_i$ of T such that every existential sentence that holds in $(\underline{B}_i, \bar{a}_i)$ also holds in $(\underline{A}_i, \bar{a}_i)$. By Lemma 7.1.6 there are an elementary extension $\underline{A}_{i+1}$ of $\underline{A}_i$ and an embedding $g_i \colon \underline{B}_i \to \underline{A}_{i+1}$ such that $g_i \circ f_i$ is the identity on $\underline{A}_i$. Then $\underline{C} := \bigcup_{i<\omega} \underline{A}_i$ equals $\lim_{i<\omega} \underline{B}_i$, and by the Tarski-Vaught elementary chain theorem (Lemma 3.1.3) $\underline{A}_0$ is an elementary substructure of $\underline{C}$. So $\underline{C}$ is a model of Φ. $\square$

7.2. Model Completeness

A theory T is *model complete* if every embedding between models of T is elementary, i.e., preserves all first-order formulas. An equivalent characterisation of model completeness is as follows.

Theorem 7.2.1. *Let T be a theory. Then the following are equivalent.*

(1) T is model complete.
(2) Every first-order formula is modulo T equivalent to an existential formula.
(3) For every embedding $e \colon \underline{A} \hookrightarrow \underline{B}$ between models $\underline{A}$ and $\underline{B}$ of T, every tuple $\bar{a}$ of elements of A, and every existential formula ϕ, if $\underline{B} \models \phi(e(\bar{a}))$ then $\underline{A} \models \phi(\bar{a})$.
(4) Every existential formula is modulo T equivalent to a universal formula;
(5) Every first-order formula is modulo T equivalent to a universal formula.

Proof. (1) $\Rightarrow$ (2). Suppose that T is model complete, and let ϕ be a first-order formula. Since T is model complete, ϕ is preserved by all embeddings between models of T. It follows from Corollary 7.1.4 that ϕ is equivalent to an existential formula.

(2) $\Rightarrow$ (3). Let e be an embedding of a model of $\underline{A}$ of T into a model $\underline{B}$ of T. Let $\bar{a}$ be a tuple of elements of $\underline{A}$, and ϕ an existential formula such that $\underline{B} \models \phi(e(\bar{a}))$. By (2), the formula $\neg\phi$ is equivalent to an existential formula. Therefore, e preserves $\neg\phi$. Since $\underline{B} \models \phi(e(\bar{a}))$ we therefore must have $\underline{A} \models \phi(\bar{a})$.

(3) $\Rightarrow$ (4): Let ϕ be a first-order formula; we have to show that $\neg\phi$ is equivalent to an existential formula. But (3) implies that $\neg\phi$ is preserved by embeddings between models of T, so the statement follows from Corollary 7.1.4.

(4) $\Rightarrow$ (5). Let ϕ be a first-order formula, written in prenex normal form $Q_1 x_1 \cdots Q_n x_n.\psi$ for ψ quantifier-free. If $n = 0$ then there is nothing to be shown. Otherwise, let $i \leq n$ be smallest so that $Q_i = \cdots = Q_n$. If $i = 1$ then ϕ is already universal, or equivalent to a universal formula by (4), and we are done. Otherwise, we distinguish two cases:

- If $Q_i = \cdots = Q_n = \exists$ then by (4) the formula $\exists x_i \ldots \exists x_n.\psi$ is equivalent modulo T to a universal formula ψ'. Then $\phi' := Q_1 x_1 \cdots Q_{i-1} x_{i-1}.\psi'$ is equivalent to ϕ.

- If $Q_i = \cdots = Q_n = \forall$ then by (4) the formula $\exists x_i \dots \exists x_n . \neg\psi$ is equivalent modulo T to a universal formula ψ'. In this case, the formula $\phi' := Q_1 x_1 \cdots Q_{i-1} x_{i-1} . \neg\psi'$ is equivalent to ϕ.

In both cases, ϕ' has fewer quantifier alternations. The claim therefore follows by induction on the number of quantifier alternations of ϕ.

(5) $\Rightarrow$ (1). Follows from the fact that existential formulas are preserved by embeddings. □

Theorem 7.2.1 shows that in particular theories with quantifier-elimination are model complete. We say that a structure $\underline{A}$ is model complete if and only if the first-order theory $\mathrm{Th}(\underline{A})$ of $\underline{A}$ is model complete.

EXAMPLE 40. All finite structures $\underline{A}$ are model complete: self-embeddings of $\underline{A}$ are automorphisms, and hence they are elementary. Every relation that is first-order definable in a finite structure also has an existential definition. On the other hand, finite structures might not have quantifier elimination. △

EXAMPLE 41. The structure $(\mathbb{Q}_0^+; <)$, where $\mathbb{Q}_0^+$ denotes the non-negative rational numbers, is not model complete, because the self-embedding $x \mapsto x + 1$ of $(\mathbb{Q}_0^+; <)$ does not preserve the formula $\phi(x) = \forall y . \neg(y < x)$ (which is satisfied only by 0). △

EXAMPLE 42. Let T be the first-order theory of $(\mathbb{Z}; \mathrm{succ})$ where succ is the binary relation $\{(x, y) \mid y = x + 1\})$. Then T does not have quantifier elimination, but is model complete. △

PROPOSITION 7.2.2. *Every model complete theory T is equivalent to a $\forall\exists$-theory.*

PROOF. Model completeness implies that any chain of models of T is an elementary chain. So unions of chains are again models by Tarski's elementary chain theorem, and the statement following from the preservation theorem of Chang-Łoś-Suszko (Theorem 7.1.8). □

Exercises.

(76) Show that $\mathrm{Th}(\mathbb{Q}; 0, 1, +, \cdot)$ is not model complete. You may use the following facts:
  - There exists a first-order formula $\phi(x)$ that defines $\mathbb{Z}$ in $(\mathbb{Q}; 0, 1, +, \cdot)$ (the formula can even be chosen to be a $\forall\exists$-formula; Poonen **[24]**).
  - Every recursively enumerable subset of $\mathbb{Z}$ has an existential definition in $(\mathbb{Z}; 0, 1, +, \cdot)$ **[23]**.
  - There are recursively enumerable subsets of $\mathbb{Z}$ whose complement is not recursively enumerable (see, e.g., **[3]**).

7.3. Model Companions

In this section we study conditions that imply that we can pass from a theory T to a model complete theory T' satisfying the same universal sentences. We write $T_\forall$ ($T_\exists$) for the set of all universal (existential) sentences that are implied by T. We say that two theories T and T' are *companions* if $T_\forall = T'_\forall$. First we observe the following consequence of Lemma 7.1.6.

COROLLARY 7.3.1. *If $T_\forall \subseteq T'_\forall$ then every model of T can be embedded into a model of T'.*

PROOF. Let $\underline{A}$ be a model of T', and let $S := \mathrm{Th}(\underline{A})_\exists$. We claim that $T \cup S$ is satisfiable. If not, then by compactness (Theorem 2.0.1) there is some finite subset $\{\phi_1, \dots, \phi_k\}$ of S such that $T \cup \{\phi_1, \dots, \phi_k\}$ is unsatisfiable. Note that $\neg\phi_1 \vee \cdots \vee \neg\phi_k$

is equivalent to a universal sentence ψ, and T implies ψ, so by assumption we have that T' implies ψ, and in particular $\underline{A} \models \psi$. We have reached a contradiction since $\underline{A} \models \phi_i$ for all $i \in \{1, \dots, k\}$. So indeed $T \cup S$ has a model $\underline{B}$. Then $\underline{B}$ satisfies the assumptions from Lemma 7.1.6, so there exists an elementary extension $\underline{C}$ of $\underline{B}$ that embeds $\underline{A}$. □

PROPOSITION 7.3.2. *Two theories T and T' are companions if and only if every model of T can be embedded into a model of T', and vice versa.*

PROOF. If T and T' are companions, then the statement follows from two symmetric applications of Corollary 7.3.1.

For the converse direction of the statement, let ϕ be a universal sentence. If ϕ is not implied by T', then there exists a model $\underline{A}$ of $T' \cup \neg\phi$. By assumption, there exists an embedding of $\underline{A}$ into a model $\underline{B}$ of T. Note that $\neg\phi$ is equivalent to an existential sentence, and hence it is preserved by the embedding, and must hold in $\underline{B}$. Hence, $\underline{B}$ cannot satisfy ϕ. This shows that every universal sentence that is implied by T is also implied by T'. The statement is symmetric in T and T', so we obtain that $T_\forall = T'_\forall$. □

DEFINITION 7.3.3. A theory T' is a *model companion* of a theory T if

- T and T' are companions, and
- T' is model complete.

EXAMPLE 43. We have seen in Example 41 that $(\mathbb{Q}_0^+; <)$ is not model complete. However, it has a model companion: the first-order theory of $(\mathbb{Q}; <)$. △

It is known that every ω-categorical theory has a model companion [**26**]. An example of a theory without model companion is the theory of groups, as we will see later (Example 45). We will now prove that the model companion of a theory T, if it exists, is unique up to logical equivalence. To prove this, we need the concepts from the next section.

7.4. Existentially Closed Structures

DEFINITION 7.4.1. Let T be a theory. A structure $\mathfrak{A}$ is called *existentially closed (e.c.) for T* (sometimes also called *existentially complete*) if there is an embedding from $\mathfrak{A}$ to a model of T, and if for any embedding h from $\mathfrak{A}$ into a model $\mathfrak{B}$ of T, any tuple $\bar{a}$ from A, and any existential formula ϕ with $\mathfrak{B} \models \phi(h(\bar{a}))$ we have $\mathfrak{A} \models \phi(\bar{a})$.

EXAMPLE 44. Every field $\underline{K}$ that is existentially closed for the theory of fields must be algebraically closed: suppose that $f(x)$ is a polynomial of degree at least one and coefficients from K. We have to show that $f(x)$ has a root in K. We can suppose that f is irreducible; then $\underline{E} := \underline{K}[x]/(f)$ is a field which extends $\underline{K}$ and contains a root of f (see Proposition B.4.9). So $\underline{E} \models \exists y. f(y) = 0$. Since $\underline{K}$ is existentially closed, $\underline{K} \models \exists y. f(y) = 0$, too, so that f has a root already in $\underline{K}$. △

We have the following lemma about existence of existentially closed models. For a theory T, we write $T_{\forall\exists}$ for the set of $\forall\exists$ sentences that are implied by T.

LEMMA 7.4.2. *Let T be a theory with a countable signature. Then any model $\underline{A}$ of T of infinite cardinality κ embeds into a structure $\mathfrak{B}$ of cardinality κ which is existentially complete for T and a model of $T_{\forall\exists}$.*

PROOF. For simplicity of presentation, we prove the statement only for the countable case $\kappa = \omega$, but the proof can easily be modified for general cardinalities κ. Set $\mathfrak{B}_0 := \mathfrak{A}$. Having constructed a countable $\mathfrak{B}_i$ let $\{(\phi_j, \bar{a}_j) \mid j \in \mathbb{N}\}$ be an enumeration of all pairs $(\phi, \bar{a})$ where $\phi(x_1, \ldots, x_n)$ is existential and $\bar{a}$ is an n-tuple from B_i. We construct a chain $(\mathfrak{B}_i^j)_{j\in\mathbb{N}}$ of countable models of T as follows.

Set $\mathfrak{B}_i^0 = \mathfrak{B}_{i-1}$, and let $j \in \mathbb{N}$. If there is a model $\underline{C}$ of T and an embedding $e\colon \mathfrak{B}_i^j \to \mathfrak{C}$ such that $\underline{C} \models \phi_j(e(\bar{a}_j))$, then by the theorem of Löwenheim-Skolem there is also a countable model $\underline{C}'$ of T and an embedding $e'\colon \mathfrak{B}_i^j \to \underline{C}'$ such that $\underline{C}' \models \phi_j(e'(\bar{a}_j))$. Identify the elements of $\underline{B}_i^j$ with elements of $\underline{C}'$ along e', and set $\mathfrak{B}_i^{j+1} := \underline{C}'$. Otherwise, if there is no such model $\underline{C}$, we set $\mathfrak{B}_i^{j+1} := \mathfrak{B}_i^j$.

Let $\underline{B}_i$ be $\lim_{j\in\mathbb{N}} \underline{B}_i^j$. Clearly, $\underline{B} := \lim_{i<\omega} \underline{B}_i$ is a countable model of $T_{\forall\exists}$. To verify that $\underline{B}$ is T-ec, first note that $\underline{B}$ embeds into a model of T by Lemma 7.1.6. Now suppose that g is an embedding from $\underline{B}$ to a model $\underline{C}$ of T, and suppose that there is a k-tuple $\bar{b}$ over B and an existential formula ϕ such that $\underline{C} \models \phi(g(\bar{b}))$. There is an $i \in \mathbb{N}$ such that $\bar{b} \in B_i^k$. By construction we have that $\underline{B}_{i+1} \models \phi(\bar{a})$. Thus, we also have that $\underline{B} \models \phi(\bar{b})$, which is what we had to show. $\square$

COROLLARY 7.4.3. *Let $\underline{K}$ be a field. Then $\underline{K}$ has an algebraically closed field extension.*

PROOF. Let T be the theory of fields. By Theorem 7.4.2, $\underline{K}$ embeds into a T-ec structure $\underline{F}$ which is a model of T. Note that the theory of fields is $\forall\exists$, so $\underline{F}$ is a field, too. The observations in Example 44 show that $\underline{F}$ is algebraically closed. $\square$

7.5. The Kaiser Hull

We will show that the model companion of T, if it exists, is precisely the theory of all existentially closed structures for T. To prove this, we need the following.

LEMMA 7.5.1. *Every theory T has a unique largest $\forall\exists$-theory as companion.*

PROOF. Suppose for contradiction that the set of $\forall\exists$-theories that are companions of T is not closed under unions. Then by compactness there are $\forall\exists$-theories S and S' and an existential sentence ϕ such that

- S and S' are companions of T
- $T \cup \{\phi\}$ has a model $\underline{A}$, and
- $S \cup S' \cup \{\phi\}$ is unsatisfiable.

By Corollary 7.3.1 there exists an embedding from $\underline{A}$ to a model $\underline{A}_0$ of S, and an embedding from $\underline{A}_0$ to a model $\underline{A}_1$ of S'. Repeating this step we construct a chain of structures $(\underline{A}_i)_{i\in\mathbb{N}}$ such that $\underline{A}_i$ is a model of S for even i and a model of S' for odd i. The union $\underline{B} := \lim_{i<\omega} \underline{A}_i$ is a model of $S \cup S'$. Since there is an embedding from $\underline{A}$ to $\underline{B}$ and $\underline{A} \models \phi$ we have that $\underline{B} \models \phi$, a contradiction. $\square$

The theory constructed in Lemma 7.5.1 will be called the *Kaiser hull* of T, denoted in the following by T^{KH}.

LEMMA 7.5.2. *T^{KH} equals the set of $\forall\exists$-sentences that hold in all T-ec structures.*

PROOF. Let T^* be the set of all $\forall\exists$-sentences satisfied by all T-ec structures. To show that $T^* \subseteq T^{\mathrm{KH}}$ it suffices to show that $(T^*)_\forall = T_\forall$. By Lemma 7.4.2, every model of $T_\forall$ embeds into a T^*-ec structure. Therefore, every universal sentence that holds in all T^*-ec structures must be in $T_\forall$, i.e., $(T^*)_\forall \subseteq T_\forall$. Conversely, every T-ec structure embeds into a model of $T_\forall$, and therefore satisfies $T_\forall$, so $T_\forall \subseteq (T^*)_\forall$.

We now show that $T^{\mathrm{KH}} \subseteq T^*$. For this, we have to show that every T-ec structure $\underline{A}$ satisfies all $\phi \in T^{\mathrm{KH}}$. Since T^{KH} is a $\forall\exists$-theory, ϕ is of the form $\forall\bar{y}.\psi(\bar{y})$ where ψ is existential. Let $\bar{a}$ be a tuple of elements of $\underline{A}$. We have to show that $\underline{A} \models \psi(\bar{a})$. Since $(T^*)_\forall = T_\forall = T^{\mathrm{KH}}_\forall$, by Corollary 7.3.1 there is an embedding e from $\underline{A}$ to a model $\underline{B}$ of $T^{\mathrm{KH}}_\forall$. Since $\underline{B} \models \phi$ we have that $\underline{B} \models \psi(e(\bar{a}))$. Since $T^{\mathrm{KH}}_\forall = T_\forall$, by Corollary 7.3.1

there is an embedding g from $\underline{B}$ to a model $\underline{C}$ of T. Since g preserves ψ we have that $\underline{C} \models \psi(g(e(\bar{a})))$. Now $\underline{A} \models \psi(\bar{a})$ since $\underline{A}$ is T-ec. We conclude that $\underline{A} \models T^{\mathrm{KH}}$. □

THEOREM 7.5.3 (Theorem 3.2.14 in [**29**]). *Let T be a theory. Then the following are equivalent.*

(1) T has a model companion;
(2) all models of the Kaiser hull of T are T-ec;
(3) the class of T-ec structures has a first-order axiomatization.

In particular, if T has a model companion T^*, then it is unique up to equivalence of theories, and $T^* = T^{\mathrm{KH}}$ is the theory of all T-ec structures.

PROOF. (1) ⇒ (2). Let U be the model companion of T. By Proposition 7.2.2, U is equivalent to a $\forall\exists$-theory. Since $U_\forall = T_\forall$ we therefore have $U \subseteq T^{\mathrm{KH}}$. So it suffices to show that every model $\underline{A}$ of U is T-ec. The structure $\underline{A}$ embeds into a model of T since $\underline{A} \models U_\forall (= T_\forall)$ by Proposition 7.3.1. Let e be an embedding from $\underline{A}$ to a model $\underline{B}$ of T, let $\bar{a}$ be a tuple from $\underline{A}$, and ϕ an existential formula with $\underline{B} \models \phi(e(\bar{a}))$. Then $\underline{B}$ has an embedding g to a model $\underline{C}$ of U. Since U is a model complete core theory, the embedding $g \circ h$ is elementary. Since g preserves existential formulas, $\underline{C} \models \phi(g(e(\bar{a})))$. Since $g \circ e$ is elementary, $\underline{A} \models \phi(\bar{a})$.

(2) ⇒ (3). Lemma 7.5.2 implies that T^{KH} is satisfied by all T-ec structures. Together with (2) this implies that the first-order theory T^{KH} axiomatises the class of all T-ec structures.

(3) ⇒ (1). Suppose that the class of T-ec structures equals the class of all models of a first-order theory U. We claim that U is the core companion of T. Every model of U is in particular a model of $T_\forall$, and every model of $T_\forall$ homomorphically maps to a model of U by Lemma 7.4.2. So we only have to verify that U is model complete. Let e be an embedding between models $\underline{A}$ and $\underline{B}$ of U, let $\bar{a}$ be a tuple from A, and let ϕ an existential formula. Since $\underline{B}$ satisfies $T_\forall$ and $\underline{A}$ is $T_\forall$-ec we have that $\phi(\bar{a})$, as desired.

The final statement of the theorem is a clear consequence of the proof above. □

EXAMPLE 45. The theory T of groups does not have a model companion. By Theorem 7.5.3 it suffices to show that the class of T-e.c. structures does not have a first-order axiomatisation. In the following, a structure that is T-e.c. is called an *existentially closed group*. We need the following group-theoretic fact, which follows from Cayley's representation theorem for groups.

Claim. $a, b \in G$ have the same order in $\underline{G}$ if and only if there is a supergroup $\underline{H}$ of $\underline{G}$ where a and b are conjugate, i.e., there exists an $h \in H$ such that $h^{-1}ah = b$.

It follows from this claim that in existentially closed groups two elements have the same order if and only if they are conjugate. Clearly, there are groups that contain elements of arbitrarily large finite order. Every group can be extended to an existentially closed group, so there is an existentially closed group $\underline{G}$ which contains elements of arbitrarily large finite order. Therefore, for each $n \in \mathbb{N}$ there are elements of order $\geq n$ which are not conjugate in $\underline{G}$. By compactness, $\underline{G}$ has an elementary extension $\underline{G}'$ which has two elements of infinite order which are not conjugate in $\underline{G}'$. Thus, $\underline{G}'$ is not existentially closed, and the class of existentially closed groups is not elementary. △

Exercises.

(77) Prove that the model companion of the theory of equivalence relations is the theory of equivalence relations with infinitely many infinite classes.
(78) Prove that the model companion of the theory of loopless undirected graphs is the theory of the countably infinite random graph.

(79) Let us call a partial order $(P; \leq)$ *tree-like* if for all $a, b, c \in P$ such that $a \leq b$ and $a \leq c$, then b and c must be comparable in $(P; \leq)$. Show that the theory of tree-like partial orders has an ω-categorical model companion.
(80) (*) Does the theory of monoids have a model companion?
(81) (*) Does the theory of rings have a model companion?
(82) (*) Does the theory of abelian groups have a model companion?

7.6. Quantifier Elimination

In Section 6.2 we studied quantifier elimination in ω-categorical theories. In this section we present characterisations of quantifier elimination for general theories.

LEMMA 7.6.1. *A first-order formula $\phi(x_1, \dots, x_n)$ is equivalent to a quantifier-free formula modulo a theory T if and only if ϕ is preserved by isomorphisms between substructures of models of T.*

PROOF. Let $\underline{A}, \underline{B}$ be models of T and $a \in A^n$. Suppose that h is an isomorphism between a substructure of $\underline{A}$ that contains a and a substructure of $\underline{B}$. Then clearly, if ϕ is quantifier-free and $\underline{A} \models \phi(a)$, then $\underline{B} \models \phi(h(a))$.

Now suppose that conversely ϕ is a first-order formula that is preserved by isomorphisms between substructures of models of T. Let Ψ be the set of all quantifier-free formulas $\psi(x_1, \dots, x_n)$ that are implied by ϕ modulo T. To show that Ψ implies ϕ modulo T, let $\underline{A} \models T$ and $a \in A^n$ be such that $\underline{A} \models \Psi(a)$. Let Θ be the set of all quantifier-free formulas θ such that a satisfies $\neg\theta$ in $\underline{A}$. Then $T \cup \{\neg\theta \mid \theta \in \Theta\} \cup \{\phi\}$ is satisfiable, because otherwise by compactness there would a be a finite subset Θ' of Θ such that $T \cup \{\neg\theta \mid \theta \in \Theta'\} \cup \{\phi\}$ is unsatisfiable. But then $\psi := \bigvee_{\theta \in \Theta'} \theta$ is a quantifier-free formula such that $T \cup \{\phi\}$ implies ψ, and hence $\psi \in \Psi$. So a satisfies ψ in $\underline{A}$, in contradiction to the assumption that $\underline{A} \models \neg\theta(a)$ for all $\theta \in \Theta$. So there exists a structure $\underline{B}$ and $b \in B^n$ which satisfy $T \cup \cup\{\neg\theta(b) \mid \theta \in \Theta\} \cup \{\phi(b)\}$.

Let $\underline{A}'$ be the smallest substructure of $\underline{A}$ that contains the entries of a and let $\underline{B}'$ be the smallest substructure of $\underline{B}$ that contains the entries of b. Since a and b satisfy the same quantifier-free formulas, the map that sends b_i to a_i can be extended to an isomorphism h between $\underline{B}'$ and $\underline{A}'$ (see Exercise 7). By assumption, h preserves ϕ, so $\underline{A} \models \phi(a)$, and indeed ψ implies ϕ modulo T. The compactness theorem implies that Ψ is equivalent to a single quantifier-free formula modulo T, as in the proof of Theorem 7.1.2. $\square$

LEMMA 7.6.2. *A theory T has quantifier elimination if and only if every formula $\exists x.\phi$ where ϕ is quantifier-free is equivalent modulo T to a quantifier-free formula.*

PROOF. For the interesting direction, let ϕ be a first-order formula. We can assume that ϕ is of the form $Q_1 x_1 \dots Q_n x_n.\psi$ for $Q_1, \dots, Q_n \in \{\exists, \forall\}$ and ψ quantifier-free. We show the statement by induction on n. For $n = 0$ there is nothing to be shown. First consider the case that $Q_n = \exists$. By assumption, $\exists x_n.\psi$ is equivalent to a quantifier-free formula ψ', so ϕ is equivalent to $Q_1 x_1 \dots Q_{n-1} x_{n-1}.\psi'$, which is in turn equivalent to a quantifier-free formula by induction.

Now consider the case that $Q_n = \forall$. Then $\exists x_n.\neg\psi$ is by assumption equivalent to a quantifier-free formula ψ', and ϕ is equivalent to $Q_1 x_1 \dots Q_{n-1} x_{n-1}.\neg\psi'$, which is equivalent to a quantifier-free formula by induction. $\square$

LEMMA 7.6.3. *A theory T has quantifier elimination if and only if for all models $\underline{B}_1$ and $\underline{B}_2$ of T with a common substructure $\underline{A}$, for all $a \in A^n$, and quantifier-free formula $\psi(y, x_1, \dots, x_n)$, if $\underline{B}_1 \models \exists y.\psi(y, a)$ then $\underline{B}_2 \models \exists y.\psi(y, a)$.*

PROOF. The forward implication is trivial. For the backward implication it suffices to show that every formula $\phi(x_1, \dots, x_n)$ of the form $\exists y.\psi$, for ψ quantifier-free, is equivalent modulo T to a quantifier-free formula (Lemma 7.6.2). By Lemma 7.6.1, it suffices to show that ϕ is preserved by isomorphisms between substructures of models of T. So let $\underline{B}_1, \underline{B}_2 \models T$ and $a \in B_1^n$, and let h be an isomorphism between a substructure that contains a and a substructure of $\underline{B}_2$. We identify the elements of the substructures along h. Clearly, a quantifier-free formula holds on a in $\underline{B}_1$ if and only if it holds on $h(a)$ in $\underline{B}_2$. The assumption then implies that $\underline{B}_1 \models \phi(a)$ if and only if $\underline{B}_2 \models \phi(h(a))$. Hence, ϕ is equivalent to a quantifier-free formula by Lemma 7.6.1. $\square$

EXAMPLE 46 (Ferrante and Rackoff [**11**], Section 3, Theorem 1). Let $\underline{L}$ be the structure with domain $\mathbb{Q}$ and the signature $\{+, 1, \leq\}$ where

- $+$ is a binary relation symbols that denotes in $\underline{L}$ the usual addition over $\mathbb{Q}$;
- 1 is a constant symbol that denotes $1 \in \mathbb{Q}$ in $\underline{L}$.
- $\leq$ is a binary relation symbol that denotes the usual linear order of the rationals.

The structure $\underline{L}$ has quantifier elimination: TODO. $\triangle$

EXAMPLE 47 (taken from [**5**]). Let $\underline{H}$ be the structure with domain $\mathbb{Q}$ and the signature $\{1, <\} \cup \{c\cdot\}_{c \in \mathbb{Q}}$ where

- 1 is a constant symbol that denotes $1 \in \mathbb{Q}$ in $\underline{L}$.
- $<$ is a binary relation symbol that denotes the usual strict linear order of the rationals.
- $\cdot c$, for $c \in \mathbb{C}$, is a unary function symbols that denotes in $\underline{L}$ the function $x \mapsto cx$ (multiplication by c).

The structure $\underline{L}$ has quantifier elimination: TODO. $\triangle$

Exercises.

(83) Find quantifier-free formulas that define the relations of $\underline{H}$ in $\underline{L}$.
(84) Show that $(\mathbb{Z}; s)$ has quantifier elimination, where $s\colon \mathbb{Z} \to \mathbb{Z}$ is given by $x \mapsto x + 1$.
(85) Show that $(\mathbb{Z}; <, s)$ has quantifier elimination.

7.7. Algebraically Closed Fields

The model companion of the theory of fields is the theory ACF of algebraically closed fields. This follows from the following important theorem.

THEOREM 7.7.1 (Tarski). ACF *has quantifier elimination.*

PROOF. We use Lemma 7.6.3. Let $\underline{K}_1, \underline{K}_2 \models$ ACF and $\underline{R}$ a common substructure; then $\underline{R}$ is a ring, and in fact an integral domain (see Appendix B.1). Let $\psi(y, x_1, \dots, x_n)$ be quantifier-free and $a \in R^n$ be such that $\underline{K}_1 \models \exists y.\psi(y, a)$. We have to show that $\underline{K}_2 \models \exists y.\psi(y, a)$. For $i \in \{1, 2\}$, let f_i be the unique embedding of the field $\underline{F}$ of fractions of $\underline{R}$ into $\underline{K}_i$ which is the identity on R (see Proposition B.1.2). By Corollary B.4.12, there exists an isomorphism g between the algebraic closure $\underline{G}_1$ of $f_1(F)$ in $\underline{K}_1$ and the algebraic closure $\underline{G}_2$ of $f_2(F)$ in $\underline{K}_2$ which extends $f_2 f_1^{-1}$.

Choose an element $b_1 \in K_1$ such that $\underline{K}_1 \models \psi(b_1, a)$. If $b_1 \in G_1$, then $\underline{K}_2 \models \psi(g(b_1), a)$. Otherwise, b_1 is transcendental, and there is an isomorphism between the field extension $\underline{G}_1(b_1)$ and the rational function field $\underline{G}_1(x)$ which fixes G_1 (under the usual identification of polynomials of degree 0 with elements of G_1) and that maps b_1 to x (see Section B.4.1).

- If $\underline{K}_2$ is a proper extension of $\underline{G}_2$, arbitrarily choose $b_2 \in \underline{K}_2 \setminus G_2$ for b_2; note that b_2 is transcendental over $\underline{G}_2$, so $\underline{G}_2(b_2)$ is isomorphic to $\underline{G}_2(x)$, too. We can then extend g to an isomorphism between $\underline{G}_1(b_1)$ and $\underline{G}_2(b_2)$ which maps b_1 to b_2 and extends $f_2 f_1^{-1}$. Hence, $\underline{K}_2 \models \psi(g(b_1), a)$.
- In the case that $\underline{K}_2 = \underline{G}_2$ we take a proper elementary extension $\underline{K}_2'$ of $\underline{K}_2$, which exists by the theorem of Löwenheim-Skolem (Theorem 3.2.1) since $\underline{K}_2$ is infinite. Similarly as above it then holds that $\underline{K}_2' \models \exists y.\psi(y, a)$, and therefore $\underline{K}_2 \models \exists y.\psi(y, a)$.

Quantifier elimination follows by Lemma 7.6.3. $\square$

The theory ACF is not complete. For any prime p let

$$\mathrm{ACF}_p := \mathrm{ACF} \cup \{\underbrace{1 + \cdots + 1}_{p \text{ times}} = 0\}$$

and for $p = 0$ let

$$\mathrm{ACF}_0 := \mathrm{ACF} \cup \{\neg \underbrace{1 + \cdots + 1}_{n \text{ times}} = 0 \mid n \in \mathbb{N}\}$$

be the theory of algebraically closed fields of characteristic p.

THEOREM 7.7.2. *For $p = 0$ or p prime, the theory ACF_p is κ-categorical for any $\kappa > \aleph_0$.*

PROOF. Let $\underline{E}$ be a model of ACF_p and let $\underline{F}$ be the smallest subfield of $\underline{E}$ (also called the *prime field*; e.g., if $p = 0$ then $\underline{F} = (\mathbb{Q}; +, *, 0, 1)$). Let $\underline{G}$ be the algebraic closure of $\underline{F}$. Then any two algebraically closed fields of the same characteristic and of the same cardinality $\kappa > \aleph_0$ have transcendence bases over $\underline{G}$ (see Section B.4.4) of cardinality κ. Any bijection between these transcendence bases induces an isomorphism of the fields. $\square$

Note that infinite fields $\underline{K}$ are never ω-categorical since there are (non-isomorphic!) countable models of $\mathrm{Th}(\underline{K})$ of transcendence degree $0, 1, 2, \ldots$.

COROLLARY 7.7.3. *For $p = 0$ or p prime, ACF_p is a complete theory.*

PROOF. Follows from Theorem 7.7.2 by Vaught's test (Theorem 3.3.2). $\square$

We also obtain the following 'cross-characteristic transfer' result.

THEOREM 7.7.4 (Lefschetz principle). *Let ϕ be a first-order sentence over the signature of rings. Then the following are equivalent.*

(1) ϕ holds in the field of complex numbers.
(2) $\mathrm{ACF}_0 \models \phi$
(3) $\mathrm{ACF}_p \models \phi$ *for all but finitely many prime numbers p.*

PROOF. (1) $\Leftrightarrow$ (2) follows from the fact that the complex numbers are algebraically closed, and that ACF_0 is complete.

(2) $\Rightarrow$ (3). Suppose that $\mathrm{ACF}_p \models \phi$ co-finitely often. We claim that $T := \mathrm{ACF}_0 \cup \{\phi\}$ is satisfiable. A finite subset T' of T will mention only finitely many inequalities $p \neq 0$, so taking q large enough, $\mathrm{ACF}_q \models T'$. By compactness, T is satisfiable. Since ACF_0 is complete, $\mathrm{ACF}_0 \models \phi$.

(3) $\Rightarrow$ (2). Suppose that infinitely often, $\mathrm{ACF}_p \not\models \phi$. Since ACF_p is complete, this means that $\mathrm{ACF}_p \models \neg\phi$ co-finitely often, and we show as above that $\mathrm{ACF}_0 \models \neg\phi$. $\square$

COROLLARY 7.7.5 (Hilbert's Nullstellensatz). *Let $\underline{K}$ be an algebraically closed field and let $f_0, \ldots, f_m \in \underline{K}[x_1, \ldots, x_n]$. Then $\underline{K} \models \exists x_1, \ldots, x_n \bigwedge_{i \in \{1,\ldots,m\}} f_i = 0$ if and only if $I = (f_1, \ldots, f_m)$ is a* proper *ideal, i.e., $1 \notin I$.*

Proof. It is clear that if $a_1, \ldots, a_n \in K$ are such that $f_1(a_1, \ldots, a_n) = \cdots = f_m(a_1, \ldots, a_n)$, and $f \in I$, then $f(a_1, \ldots, a_n) = 0$, so $1 \notin I$.

For the converse, an application of Zorn's lemma to the set $\mathcal{S}$ of proper ideals J in $\underline{K}[x_1, \ldots, x_n]$ containing I gives a maximal ideal $J \in \mathcal{S}$. Since J is maximal, $\underline{L} := \underline{K}[x_1, \ldots, x_n]/J$ is a field (Proposition B.2.7). We identify $\underline{K}$ with a subfield of $\underline{L}$ by identifying $a \in K$ with $a + J$. For each $i \leq k$ we have

$$f_i(x_1 + J, \ldots, x_n + J) = f_i(x_1, \ldots, x_n) + J = J$$

since $f_i(x_1, \ldots, x_n) \in I \subseteq J$. Hence $\exists x_1, \ldots, x_n \bigwedge_{i \in \{1,\ldots,n\}} f_i = 0$ holds in $\underline{L}$, and thus also in its algebraic closure. By quantifier-elimination for ACF, this is also true for the substructure $\underline{K} \models$ ACF. $\square$

Exercises

(86) Find a quantifier-free formula which is equivalent to $\exists x(ax^2 + bx + c = 0)$ over $(\mathbb{C}; +, *)$.

(87) Find a quantifier-free formula which is equivalent to the formula of the previous exercise over $(\mathbb{R}; +, *)$.

(88) Write down a first-order formula $\phi(a, b, c, d)$ in the language of fields which states that the matrix $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ is invertible. Find a quantifier-free formula which is equivalent to ϕ for any field $\underline{F}$.

(89) (Exercise 2.3.3 in [**29**]) Show that a $\forall\exists$-sentence which holds in all finite fields is true in all algebraically closed fields.

(90) Show that in every algebraically closed field $\underline{K}$, every injective polynomial map of a definable subset of K^n to K is surjective.

CHAPTER 8

Stability

Stability theory grew out of the attempt to understand the possible spectra of first-order theories (Section 1.3). Today, stability theory might be thought of as a way to classify definable sets in a structure whenever this is possible. The starting point is that instead of counting the number of models of a first-order theory T, we *count the number of types* in models of T over parameter sets of some cardinality. This chapter presents some of the basic and fundamental facts of stability theory, and is not yet finished.

DEFINITION 8.0.1. Let κ be an infinite cardinal. A theory T is *κ-stable* if for every $\underline{M} \models T$ and for every $A \subseteq M$, if $|A| \leq \kappa$ then $|S_1^{\underline{M}}(A)| \leq \kappa$. If T is κ-stable for some κ, then T is called *stable*, and it is called *unstable* otherwise. A structure is called *κ-stable* if its first-order theory is κ-stable.

EXAMPLE 48. Let E be an equivalence relation on $\mathbb{N}$ with infinitely many infinite classes. Then the structure $(\mathbb{N}; E)$ is stable. △

EXAMPLE 49. Let p be a prime or 0. Algebraically closed fields of characteristic p are κ-stable for all κ. Recall that ACF_p is κ-categorical for all $\kappa > \omega$ (Theorem 7.7.2). We show in Theorem 8.3.2 below that uncountable categoricity implies ω-stability. Then we show that ω-stability implies κ-stability for all κ in Proposition 8.3.3. Alternatively, it is easy to prove κ-stability directly, using quantifier elimination in ACF (Theorem 7.7.1). △

The following lemma shows that it does not matter whether we define stability with respect to 1-types or with respect to n-types (similarly as with saturation, Exercise 54).

LEMMA 8.0.2. *Let κ be an infinite cardinal and let T be κ-stable. Then T is* κ-stable for n-types, *for every $n \in \mathbb{N}$, i.e., for every model $\underline{M}$ of T and $A \subseteq M$ of cardinality at most κ we have $|S_n^{\underline{M}}(A)| \leq \kappa$.*

PROOF. Our proof is by induction on n. For $n = 1$ the statement holds by assumption, so suppose that $n > 1$. Let $\underline{M}$ be a model of T and $A \subseteq M$ of cardinality at most κ. By Lemma 5.2.3, $\underline{M}$ has an elementary extension $\underline{N}$ that realises all 1-types over A. For an n-type p of $\underline{M}$ over A, let p' be the 1-type that contains all sentences of the form $\exists x_2, \dots, x_n.\phi$ for $\phi \in p$. By assumption, there are only countably many 1-types of $\underline{M}$ over A. Therefore, it suffices to show that every n-type p, the set $Q\colon \{q \in S_n^{\underline{N}}(A) \mid q' = p'\}$ has cardinality at most κ, because $\omega \cdot \kappa = \kappa$.

By assumption, p' is realised by an element $b \in N$ in $\underline{N}$. Let p'' be the $(n-1)$-type obtained from p by replacing x_1 by the constant b. Then the $(n-1)$-types of $\underline{N}$ over $A \cup \{b\}$ are in bijective correspondence to the set of all $q \in S_n^{\underline{M}}(A)$ such that $q' = p'$. By inductive assumption, there are at most κ many $(n-1)$-types of $\underline{N}$ over $A \cup \{b\}$. □

Stability has many equivalent characterisations; one is based on the *binary tree property* and another one on the *order property*; they will be presented in Section 8.1.

Another equivalent characterisation is based on indiscernible sequences, a concept for convenient use of Ramsey's theorem from combinatorics (which is recalled in Appendix C); indiscernible sequences are the topic of Section 8.2. We continue with the strongest form of stability, namely ω-stability, in Section 8.3, and prove that every uncountably categorical theory is ω-stable.

Unstable theories have the independence property or the strict order property (Section 8.4). Many of the good properties of stable theories also hold for theories that do not have the independence property; such theories are called NIP and they are currently intensively studied in model theory [**28**]. They also found applications in structural graph theory and theoretical computer science (see [**6**] and the references therein). We present an interesting connection between NIP theories and the concept of VC dimension from computational learning theory (Section 9.2).

For further reading on stability theory, I recommend [**9**].

8.1. The Order Property

Stability has many equivalent characterisations, some of which will be presented in this section. We start with the observation that if a theory is unstable, then this is witnessed by a single formula (Proposition 8.1.2). Such witnessing formulas have many equivalent characterisations (Theorem 8.1.5). Throughout this section, T denotes a complete theory over the signature τ.

The following definitions about types were essentially already given in Chapter 5. Let $\underline{M}$ be a model of T and $B \subseteq M$. For a τ-formula $\phi(\bar{x}, \bar{y})$, a *(complete) ϕ-type in $\underline{M}$ over B* is a maximal set p of formulas of the form $\phi(\bar{x}, \bar{b})$ or of the form $\neg\phi(\bar{x}, \bar{b})$ for tuples $\bar{b}$ of elements of B such that p is satisfiable over some elementary extension of $\underline{M}$. Define $S^{\underline{M}}_{\phi}(B)$ to be the set of all ϕ-types in $\underline{M}$ over B.

Definition 8.1.1. A formula ϕ is called *stable* in T if for every infinite model $\underline{B}$ of T we have $|S^{\underline{B}}_{\phi}(B)| \leq |B|$, and *unstable* otherwise.

The terminology is motivated by the following corollary.

Proposition 8.1.2. *T is unstable if and only if there exists a τ-formula ϕ which is unstable in T.*

Proof. Suppose that $\phi(\bar{x}, \bar{y})$ is unstable in T, that is, there exists an infinite model $\underline{B}$ of T such that $|S^{\underline{B}}_{\phi}(B)| > |B|$. It follows that T is not $|B|$-stable, and hence unstable.

Now suppose that every τ-formula $\phi(\bar{x}, \bar{y})$ is stable. Let $\underline{M}$ be a model of T and let $B \subseteq M$ be of some infinite cardinality κ. Then $|S^{\underline{M}}_{\phi}(B)| \leq |B|$. Since every $p \in S^{\underline{M}}_{1}(B)$ is determined by $\{q \in S^{\underline{M}}_{\phi}(B) \mid \phi \text{ a } \tau\text{-formula}, q \subseteq p\}$ we have that $|S^{\underline{M}}_{1}(B)| \leq \kappa^{\omega} = \kappa$ and hence T is stable. □

How can we show that a formula ϕ is unstable in T? We have to find for every model of T and every cardinal κ strictly more than κ ϕ-types if we are allowed κ many parameters. The idea how to find so many types is to construct a big tree where the infinite branches correspond to ϕ-types.

If A and B are sets, we write A^B for the set of all functions from B to A. If $w \in A^n$ then we define $|w| := n$ and treat w as a *word* whenever this is convenient; e.g., 0100 denotes a word $w \in \{0,1\}^4$. For an ordinal μ we write $A^{<\mu}$ for $\bigcup_{n<\mu} A^n$. A *binary tree over A* is a family of elements of A of the form $(a_s)_{s \in \{0,1\}^{<\omega}}$.

Definition 8.1.3. A formula $\phi(\bar{x}, \bar{y})$ has the *binary tree property* in T if there exists a model $\underline{M}$ of T and a binary tree $(\bar{b}_w)_{w \in \{0,1\}^{<\omega}}$ of tuples of elements of M

such that for every $w \in \{0,1\}^{<\omega}$

$$\{\phi(\bar{x}, \bar{b}_{w|_n}) \mid n < \omega, w(n) = 0\} \cup \{\neg\phi(\bar{x}, \bar{b}_{w|_n}) \mid n < \omega, w(n) = 1\}.$$

is satisfiable in $\underline{M}$.

EXAMPLE 50. If $\underline{M} = (\mathbb{Q}; <)$ then the formula $\phi(x, y_1, y_2)$ given by $y_1 < x < y_2$ has the binary tree property in $\mathrm{Th}(\underline{M})$. $\triangle$

The next property, called the *order property*, is often more convenient to verify than the binary tree property, but in fact implies the binary tree property.

DEFINITION 8.1.4. Let $\phi(\bar{x}, \bar{y})$ be a first-order formula. Then ϕ has the *order property* in T if there exists a model $\underline{M}$ of T and sequences $(\bar{a}_i)_{i\in\omega}$ and $(\bar{b}_i)_{i\in\omega}$ of tuples of elements of M such that

$$\underline{M} \models \phi(\bar{a}_i, \bar{b}_j) \Leftrightarrow i < j.$$

EXAMPLE 51. If $\underline{M} = (\mathbb{Q}; <)$ then the formula $\phi(x, y)$ given by $x < y$ has the order property in $\mathrm{Th}(\underline{M})$. $\triangle$

EXAMPLE 52. If $\underline{M} = (\mathbb{V}; E)$ is the Rado graph then the formula $E(x, y)$ has the order property in $\mathrm{Th}(\underline{M})$: pick sequences $(a_n)_{n\in\omega}$ and $(b_n)_{n\in\omega}$ of elements of $\mathbb{V}$ such that $E(a_i, b_j)$ if and only if $i < j$; such elements exist by the universality of the Rado graph. $\triangle$

THEOREM 8.1.5. *Let T be a theory and let $\phi(\bar{x}; \bar{y})$ be a τ-formula. The following are equivalent.*

1. *ϕ is stable in T, i.e., for every infinite model $\underline{B}$ of T we have $|S^{B}_{\phi}(B)| \leq |B|$.*
2. *For every model $\underline{M}$ of T there exists an infinite cardinal λ such that for every $B \subseteq M$ with $|B| \leq \lambda$ we have $|S^{M}_{\phi}(B)| \leq |B|$.*
3. *$\phi(\bar{x}; \bar{y})$ does not have the binary tree property in T.*
4. *$\phi(\bar{x}; \bar{y})$ does not have the order property in T.*

PROOF. The implication from (1) to (2) is immediate.

(2) $\Rightarrow$ (3): We prove the contraposition. Let $\phi(\bar{x}, \bar{y})$ be a formula with the binary tree property in T and let λ be an infinite cardinal. Choose μ minimal such that $\lambda < 2^{\mu}$. Then $|\{0,1\}^{<\mu}| \leq \lambda$. The binary tree property and compactness imply that T has a model $\underline{M}$ and a sequence $(\bar{b}_\rho)_{\rho\in\{0,1\}^{<\mu}}$ of tuples of elements of M such that for every $\sigma \in \{0,1\}^{<\mu}$

$$p_\sigma \coloneqq \{\phi(\bar{x}, \bar{b}_{\sigma|_\alpha}) \mid \alpha < \mu, \sigma(\alpha) = 0\} \cup \{\neg\phi(\bar{x}, \bar{b}_{\sigma|_\alpha}) \mid \alpha \leq \mu, \sigma(\alpha) = 1\}$$

is satisfiable in $\underline{M}$. Complete every p_σ to an element of $S^{M}_{\phi}(B)$ where B is the set of all entries of all tuples in $(\bar{b}_\rho)_{w\in\{0,1\}^{<\mu}}$. Since all the q_σ are pairwise distinct, we have

$$|B| \leq \lambda < 2^{\mu} \leq S^{M}_{\phi}(B).$$

(3) $\Rightarrow$ (4): We again prove the contraposition. Let $\phi(\bar{x}, \bar{y})$ be a formula with the order property in T, so there is a model $\underline{M}$ and sequences $(\bar{a})_{i\in\omega}$ and $(\bar{b})_{i\in\omega}$ such that $\underline{M} \models \phi(\bar{a}_i, \bar{b}_j) \Leftrightarrow i < j$. Let $<$ be the *infix ordering* on $I \coloneqq \{0,1\}^{\omega} \cup \{0,1\}^{<\omega}$, which is defined for $u, v \in I$ by

$$\begin{array}{llll} u < v & \text{if} & v = u1x & \text{for some } x \in I \\ & & \text{or } u = w0x \text{ and } v = w1y & \text{for some } w \in \{0,1\}^{<\omega} \text{ and } x, y \in I. \end{array}$$

It suffices to prove that there are sequences $(\bar{a}'_u)_{u\in I}$ and $(\bar{b}'_u)_{u\in I}$ of tuples of M such that $\underline{M} \models \phi(\bar{a}'_u, \bar{b}'_v)$ if and only if $u < v$. The existence of such sequences follows from the compactness theorem and the existence of $(\bar{a})_{i\in\omega}$ and $(\bar{b})_{i\in\omega}$.

(4) $\Rightarrow$ (1). Again we prove the contraposition. Let $\underline{M}$ be an infinite model of T such that $|S_\phi^{\underline{M}}(M)| > |M|$. Let $\underline{M}'$ be an elementary extension of $\underline{M}$ that realises all the ϕ-types in $S_\phi^{\underline{M}}(M)$ (Lemma 5.2.3). Let $|\bar{y}| = n$. Note that every $p \in S_\phi^{\underline{M}}(M)$ is uniquely given by the set $S_p \subseteq M^n$ of all $\bar{b} \in M^n$ such that $\phi(\bar{x}, \bar{b}) \in p$. Let $\mathcal{S} := \{S_p \mid p \in S_\phi^{\underline{M}}(M)\}$. Applying the Erdős-Makkai theorem (Theorem C.0.5) to $A := M^n$ and $\mathcal{S}$ we obtain a sequence $(a_i)_{i\in\omega}$ of elements of A and a sequence $(S_i)_{i\in\omega}$ of elements of $\mathcal{S}$ such that either for all $i, j \in \omega$ we have $a_i \in S_j \Leftrightarrow i < j$, or for all $i, j \in \omega$ we have $a_i \in S_j \Leftrightarrow j < i$. In the first case, $\phi(\bar{x}, \bar{y})$ has the order property in T. Note that the definition of the order property is symmetric, so also in the second case ϕ has the order property in T. $\square$

Stability has another elegant characterisation, based on item (2) of Theorem 8.1.5.

Theorem 8.1.6. *T is unstable if and only if there exists a τ-formula $\phi(\bar{x}, \bar{y})$, a model $\underline{M}$ of T, and a sequence $(\bar{c}_i)_{i\in\omega}$ such that $\underline{M} \models \phi(\bar{c}_i, \bar{c}_j) \Leftrightarrow i < j$.*

Proof. If T is unstable, then by Proposition 8.1.2 there exists a formula $\phi(\bar{x}, \bar{y})$ which is unstable in T. By Theorem 8.1.5, ϕ has the order property in T, so there exists a model $\underline{M}$ of T and sequences $(\bar{a}_i)_{i\in\omega}$ and $(\bar{b}_i)_{i\in\omega}$ of tuples from M such that $\underline{M} \models \phi(\bar{a}_i, \bar{b}_j)$ if and only if $i < j$. Let $\psi(\bar{x}_1, \bar{x}_2, \bar{y}_1, \bar{y}_2)$ be the formula $\phi(\bar{x}_1, \bar{y}_2)$. For $i \in \omega$, define $c_i := (a_i, b_i)$. Then $\underline{M} \models \psi(\bar{c}_i, \bar{c}_j)$ if and only if $\underline{M} \models \phi(\bar{a}_i, \bar{b}_j)$ if and only if $i < j$. The backwards direction is immediate, since we may take $(b_i)_{i\in\omega} = (a_i)_{i\in\omega} = (c_i)_{i\in\omega}$. $\square$

8.2. Indiscernible Sequences

Let $\underline{A}$ be a τ-structure. Let I be a linearly ordered set and $(a_i)_{i\in I}$ is a sequence of elements of A^n for some $n \in \mathbb{N}$. Then $(a_i)_{i\in I}$ is called an *order indiscernible sequence (in $\underline{A}$)*[1] if for every τ-formula $\phi(\bar{x}_1, \dots, \bar{x}_k)$ and all elements $i_1, \dots, i_k, j_1, \dots, j_k \in I$ with $i_1 < \dots < i_k$ and $j_1 < \dots < j_k$ we have that $\underline{A} \models \phi(\bar{a}_{i_1}, \dots, \bar{a}_{i_k}) \Leftrightarrow \phi(\bar{b}_{i_1}, \dots, \bar{b}_{i_k})$. Note that if $\bar{a}_i = \bar{a}_j$ for distinct elements $i, j \in I$, then $a_i = a_j$ for all elements $i, j \in I$. This is why we usually assume that the $\bar{a}_i$ are pairwise distinct.

Example 53. Any strictly increasing or strictly decreasing sequence of elements of $\underline{A} := (\mathbb{Q}; <)$ is order indiscernible. $\triangle$

Example 54. Let $\underline{A}$ be a vector space and let X a basis of $\underline{A}$. Then any sequence $(a_i)_{i\in I}$ of elements of X is order indiscernible in $\underline{A}$. To see this, let $\phi(x_1, \dots, x_k)$ be a τ-formula and let $i_1, \dots, i_k, j_1, \dots, j_k \in I$ be such that $i_1 < \dots < i_k$ and $j_1 < \dots < j_k$. Since X is a basis of $\underline{A}$, there is an automorphism of $\underline{A}$ which takes $\bar{a} := (a_{i_1}, \dots, a_{i_k})$ to $\bar{b} := (b_{j_1}, \dots, b_{j_k})$. Hence, $\underline{A} \models \phi(\bar{a})$ if and only if $\underline{A} \models \phi(\bar{b})$. $\triangle$

Remark 8.2.1. The order indiscernible sequences in the previous example are indiscernible in a very strong sense. We say that $(a_i)_{i\in I}$ is *totally indiscernible (in $\underline{A}$)* if $(a_i)_{i\in I}$ is order indiscernible for every linear ordering of I. Example 53 shows an order indiscernible set which is not totally indiscernible.

Definition 8.2.2. Let $\underline{A}$ be a τ-structure. Let I be an infinite linearly ordered set and let $(\bar{a}_i)_{i\in I}$ be a sequence of elements of A^n for some $n \in \mathbb{N}$. Then the *Ehrenfeucht-Mostowski type $EM((\bar{a}_i)_{i\in I})$* is the set of τ-formulas $\phi(\bar{x}_1, \dots, \bar{x}_k)$, $k \in \mathbb{N}$, such that $\phi(\bar{a}_{i_1}, \dots, \bar{a}_{i_k})$ holds for all $i_1, \dots, i_k \in I$ with $i_1 < \dots < i_k$ in some elementary extension of $\underline{A}$.

---

[1]Order indiscernible sequences are sometimes just called *indescernible.*

LEMMA 8.2.3. *Let $\underline{A}$ be an infinite τ-structure. Let I be an infinite linearly ordered sets and let $(\bar{a}_i)_{i \in I}$ be a sequence of elements of A^n for some $n \in \mathbb{N}$. Then for any linearly ordered set J there exists a model $\underline{B}$ of* $\mathrm{Th}(\underline{A})$ *and an order indiscernible sequence $(\bar{b}_i)_{i \in J}$ of elements of B^n with the same Ehrenfeucht-Mostowski type as $(\bar{a}_i)_{i \in I}$.*

PROOF. By Ramsey's theorem and compactness. For every $j \in J$ and $i \in \{1, \dots, n\}$ let c_j^i be a new constant symbol; we write $\bar{c}_j$ for $(c_j^1, \dots, c_j^n)$. Let $C := \{\bar{c}_j \mid j \in J\}$. Define

$$S := \{\phi(\bar{c}_{i_1}, \dots, \bar{c}_{i_k}) \mid k \in \mathbb{N}, \phi(\bar{x}_1, \dots, \bar{x}_k) \in \mathrm{EM}((\bar{a}_i)_{i \in I}), i_1, \dots, i_k \in J, i_1 < \cdots < i_k\}$$
$$T := \{\phi(\bar{c}_{i_1}, \dots, \bar{c}_{i_k}) \Leftrightarrow \phi(\bar{c}_{j_1}, \dots, \bar{c}_{j_k}) \mid k \in \mathbb{N}, i_1, \dots, i_k, j_1, \dots, j_k \in J,$$
$$i_1 < \cdots < i_k, j_1 < \cdots < j_k, \phi \text{ a } \tau\text{-formula}\}$$

We have to show that $\mathrm{Th}(\underline{A}) \cup S \cup T$ is satisfiable. By compactness, it is enough to show that for every $k \in \mathbb{N}$, every finite set Δ of τ-formulas $\phi(\bar{x}_1, \dots, \bar{x}_k)$, and every finite subset $C_0 \subseteq C$ the theory $\mathrm{Th}(\underline{A}) \cup S' \cup T'$ is satisfiable where

$$S' := \{\phi(\bar{c}_{i_1}, \dots, \bar{c}_{i_k}) \in S \mid \bar{c}_{i_1}, \dots, \bar{c}_{i_k} \in C_0\}$$
$$T' := \{\phi(\bar{c}_{i_1}, \dots, \bar{c}_{i_k}) \Leftrightarrow \phi(\bar{c}_{j_1}, \dots, \bar{c}_{j_k}) \in T \mid \phi \in \Delta, \bar{c}_{i_1}, \dots, \bar{c}_{i_k}, \bar{c}_{j_1}, \dots, \bar{c}_{j_k} \in C_0\}.$$

Define an equivalence relation $\sim$ on the set D of k-tuples $(\bar{c}_{i_1}, \dots, \bar{c}_{i_k}) \in C^k$ such that $i_1 < \cdots < i_k$: for $\bar{u}, \bar{v} \in D$, define $\bar{u} \sim \bar{v}$ if for all $\phi \in \Delta$

$$\underline{A}' \models \phi(\bar{u}) \Leftrightarrow \phi(\bar{v})$$

holds in some elementary extension $\underline{A}'$ of $\underline{A}$. Since this equivalence relation has at most 2^{Δ} classes, by Ramsey's theorem there exists an infinite subset $C_1 \subseteq C$ such that all tuples in $C_1^k \cap D$ lie in the same class. Let $\{\bar{c}_{p_1}, \dots, \bar{c}_{p_s}\} = C_0$ with $p_1 < \cdots < p_s$. Choose $\bar{c}_{q_1}, \dots, \bar{c}_{q_s} \in C_1$ with $q_1 < \cdots < q_s$. Then the expansion of the τ-reduct of $\underline{A}'$ by interpreting $\bar{c}_{p_i}$ by $\bar{c}_{q_i}$ is a model of $\mathrm{Th}(\underline{A}) \cup S' \cup T'$. $\square$

COROLLARY 8.2.4. *Let T be a τ-theory with infinite models. Then for any linearly ordered set I there exists a model $\underline{A}$ of T with an order indiscernible sequence $(a_i)_{i \in I}$ of distinct elements of A.*

PROOF. By the upwards Löwenheim-Skolem theorem (Theorem 3.2.1) there exists a model of T with a sequence of distinct elements indexed by I. Apply Lemma 8.2.3 to this sequence. $\square$

Exercises

(91) Let T be the theory ACF_0 of algebraically closed fields of characteristic 0. Find an order indiscernible sequence $(a_i)_{i \in \omega}$ of distinct elements in a model of T.

Indiscernible sequences can also be used to characterise stability.

THEOREM 8.2.5. *T is stable if and only if every order indiscernible sequence in every model $\underline{M}$ of T is totally indiscernible.*

PROOF. If T is unstable, then by Theorem 8.1.6 there exists a τ-formula $\phi(\bar{x}, \bar{y})$, a model $\underline{M}$ of T, and a sequence $(\bar{a}_i)_{i \in \omega}$ of tuples of elements of M such that for all $i, j \in \omega$ we have $\underline{M} \models \phi(\bar{a}_i, \bar{a}_j) \Leftrightarrow i < j$. Then $(\bar{a}_i)_{i \in \omega}$ is not totally indiscernible.

Conversely, suppose that T has a model $\underline{A}$ with an order indiscernible sequence $(\bar{a}_i)_{i \in \omega}$ which is not totally indiscernible. By Lemma 8.2.3 T has a model $\underline{B}$ with an order indiscernible sequence $(\bar{b}_i)_{i \in \mathbb{Q}}$ of the same Ehrenfeucht-Mostowski type as $(a_i)_{i \in \omega}$; note that $(\bar{b}_i)_{i \in \mathbb{Q}}$ is in particular not totally indiscernible. This means that there exists a formula $\phi(\bar{x}_1, \dots, \bar{x}_k)$, some indices $i_1, \dots, i_k \in \omega$ with $i_1 < \dots < i_k$ and some permutation σ of $\{1, \dots, k\}$ such that $\underline{B} \models \phi(\bar{b}_{i_1}, \dots, \bar{b}_{i_k}) \wedge \neg\phi(\bar{b}_{\sigma(i_1)}, \dots, \bar{b}_{\sigma(i_k)})$. Any permutation on $\{1, \dots, k\}$ is a product of transpositions of consecutive elements of $\{1, \dots, k\}$, and hence we may assume that σ is a transposition that exchanges $j, j+1 \in \{1, \dots, k\}$ and fixes all elements of $\{1, \dots, k\}$. Then the formula $\phi(x_1, x_2)$ given by $\phi(\bar{a}_1, \dots, \bar{a}_j, x_1, x_2, \bar{a}_{j+2}, \dots, \bar{a}_k)$ orders the infinite sequence $(\bar{a}_i)_{\{i \in I \mid i_j < i < i_{j+1}\}}$. Hence, Theorem 8.1.6 implies that T is unstable. □

LEMMA 8.2.6. *Let τ be a countable signature and let $\underline{M}$ be a τ-structure generated by a well-ordered set I of indiscernibles. Then $\underline{M}$ realises only countably many types over every countable subset of $\underline{M}$.*

PROOF. First consider the simpler case where the signature is relational. Let $S \subseteq M$ be countable. Suppose that a realises $p \in S_1^{\underline{M}}(S)$. Then p is determined by the position of a with respect to the well-ordered subset S, and there are only countably many possibilities. The general case with function symbols can be shown similarly. □

We conclude this section by showing that every countable theory has a model that realises only countably many types over every countable subset.

DEFINITION 8.2.7. Let τ be a signature. A *Skolem theory for τ* is a theory S in a bigger signature σ such that

(1) $|\sigma| \leq \max(|\tau|, \omega)$.
(2) S is universal.
(3) Every τ-structure can be expanded to a model of S.
(4) S has quantifier elimination.

THEOREM 8.2.8. *For every signature τ there exists a Skolem theory S.*

PROOF. We define an ascending sequence of signatures $\tau = \sigma_0 \subseteq \sigma_1 \subseteq \cdots$ by introducing for every quantifier-free $\{\sigma_i\}$-formula $\phi(x_1, \dots, x_n, y)$ a new n-ary *Skolem function* f_ϕ and defining σ_{i+1} as the union of σ_i and the set of these function symbols. Define $\sigma := \bigcup_i \sigma_i$. Now define

$$S := \{\forall \bar{x}(\exists y \phi(\bar{x}, y) \Rightarrow \phi(\bar{x}, f_\phi(\bar{x})) \mid \phi(\bar{x}, y) \text{ quantifier-free } \sigma\text{-formula}\}.$$

The first two items are clearly satisfied. For the third item, let $\underline{A}$ be a τ-structure. To define the expansion $\underline{B} \models S$ of $\underline{A}$, suppose $f \in \sigma$ is n-ary, and let $a \in A^n$. If $\underline{A} \models \exists y \phi(a, y)$ for ϕ quantifier-free, then choose for $f^{\underline{B}}(a)$ any $b \in A$ such that $\underline{A} \models \phi(a, b)$. Clearly, $\underline{B} \models S$. We leave the proof that S has quantifier elimination as an exercise. □

LEMMA 8.2.9. *Let T be a countable τ-theory with an infinite model and let κ be an infinite cardinal. Then T has a model of cardinality κ which realises only countably many types over every countable subset.*

PROOF. Let τ be the signature of T, and consider the theory $T^* := T \cup S$ where S is a Skolem theory for τ with signature σ (see Theorem 8.2.8). Then T^* is countable, has an infinite model and quantifier elimination. Since S is universal and has quantifier-elimination, T^* is equivalent to a universal theory. By Lemma 8.2.3, T^* has a model $\underline{N}^*$ with a set of order indiscernibles $I = \{a_i \mid i \in \kappa\}$. Then $\underline{M}^* := \underline{N}^*[I]$ is a model of T^* of cardinality κ. Since T^* has quantifier elimination, $\underline{M}^*$ is an elementary substructure of $\underline{N}^*$ and I is order indiscernible in $\underline{M}^*$. By Lemma 8.2.6 the structure $\underline{M}^*$ realises only countably many types over every countable subset of M, and the same is true for the τ-reduct $\underline{M}$ of $\underline{M}^*$. □

8.3. ω-stable Theories

An example of an ω-stable structure is $(\mathbb{Q}; =)$. Together with the following theorem, it provides many examples of (ω-categorical) ω-stable structures.

PROPOSITION 8.3.1. *Let $\underline{A}$ be an ω-categorical ω-stable structure. Then every structure $\underline{B}$ with a first-order interpretation in $\underline{A}$ is ω-stable.*

PROOF. Recall that since $\underline{B}$ has a first-order interpretation in an ω-categorical structure, it is ω-categorical as well (Lemma 6.3.2). Let n be the dimension of the interpretation of $\underline{B}$ in $\underline{A}$. Suppose that $\underline{M}$ is a model of $\mathrm{Th}(\underline{B})$ and $S \subseteq M$ is countable such that $S_1^{\underline{M}}(S)$ is uncountable. By the Theorem of Löwenheim-Skolem downwards (Theorem 3.2.1) $\underline{M}$ has a countable elementary substructure $\underline{C}$ that contains S. By the ω-categoricity of $\underline{B}$ the structure $\underline{C}$ is isomorphic to $\underline{B}$, so we may identity $\underline{B}$ and $\underline{C}$. For each $b \in S$, fix an n-tuple of elements of A that is mapped to b by the interpretation; let T be the countable set of entries of all these tuples. Let $c, d \in B$ and let c' and d' be n-tuples of elements of A that are mapped to c and to d by the interpretation. If $c, d \in B$ have distinct 1-types of $\underline{A}$ over S, then c' and d' have distinct n-types over T. Lemma 8.0.2 then implies that $\underline{A}$ also has uncountably many 1-types over T. $\square$

Many more examples of ω-stable theories come from the following theorem.

THEOREM 8.3.2. *A countable theory T which is categorical in an uncountable cardinal κ is ω-stable.*

PROOF. Let $\underline{M}$ be a model and $A \subseteq M$ countable with $S_1^{\underline{M}}(A)$ uncountable. Let $(b_i)_{i \in I}$ be a sequence of $\aleph_1$ many elements with pairwise distinct types over A. We assume that all types over A are realised in $\underline{M}$: if not, we replace $\underline{M}$ by an elementary extension of $\underline{M}$ where all types over A are realised. By Löwenheim-Skolem downwards (Theorem 3.2.1) there is an elementary substructure $\underline{M}_0$ of $\underline{M}$ of cardinality $\aleph_1$ which contains A and all b_i. By Löwenheim-Skolem upwards (Theorem 3.2.1) there is an elementary extension $\underline{N}$ of $\underline{M}_0$ of cardinality κ. The model $\underline{N}$ realises uncountably many types over A. By Lemma 8.2.9, the theory T has another model in which only countably many types over A are realised. So T is not κ-categorical. $\square$

PROPOSITION 8.3.3. *If T is ω-stable, then T is κ-stable for all κ.*

PROOF. Suppose that T is not κ-stable for some κ, i.e., there is a model $\underline{M}$ and $A \subseteq M$ such that $|A| \leq \kappa$ and $|S_1^{\underline{M}}(A)| > \kappa$. Recall our notation from Section 5.1: for a formula $\phi(x)$ in the language of $\underline{M}_A$ we write $[\phi]$ for the set of all $p \in S_1^{\underline{M}}(A)$ that contain ϕ. By assumption, $|[x = x]| > \kappa$. Suppose now that $|[\phi]| > \kappa$. Note that every formula ψ such that $|[\psi]| \leq \kappa$ belongs to at most κ types in $[\phi]$, and that there are at most κ formulas. Hence, there must be distinct types p, q such that $|[\psi]| > \kappa$ for all formulas ψ in p and in q. Let $\psi \in p \setminus q$. We then decompose $[\phi] = [\phi \wedge \psi] \cup [\phi \wedge \neg\psi]$. Inductively, we can now construct an infinite complete binary tree of formulas. Only a countably infinite subset B of constants from A are used in these formulas, but $S_1^{\underline{M}}(B) = 2^\omega$, showing that T is not ω-stable. $\square$

Exercises.

(91) Show that a countably infinite vector space over a finite field (see Exercise 2) is ω-stable.

(92) Show that every stable structure which is homogeneous in a finite relational language is ω-stable.

Every type is given by

But then over countably many parameters

8.4. The Classification Picture

The failure of stability can occur in two different ways.

Definition 8.4.1 (The strict order property). Let T be a theory. A formula formula $\phi(\bar{x}, y_1, \dots, y_m)$ has the *strict order property (SOP)* in T if there exists $\underline{A} \models T$ and a sequence $(b_i)_{i \in \omega}$ of elements of A^m such that $\phi(\bar{x}, b_i)^{\underline{A}} \subsetneq \phi(\bar{x}, b_j)^{\underline{A}}$ for all $i < j$. We say that T has the SOP if there exists a formula that has the SOP in T.

Example 55. A typical theory with the SOP is the theory of $(\mathbb{Q}; <)$. △

Proposition 8.4.2. *If a theory T has the SOP, then it is unstable.*

Proof. Let $\phi(\bar{x}, y_1, \dots, y_m)$ be a formula, let $\underline{A}$ be a model of T, and let $(b_i)_{i \in \omega}$ be a sequence of elements of A^m that witness that T has the SOP. Let $\psi(\bar{x}, \bar{x}')$ be the formula $\neg \exists \bar{y}.\, \phi(\bar{y}, \bar{x}) \wedge \exists \bar{y}.\, \phi(\bar{y}, \bar{x})$. Then $\underline{A} \models \psi(b_i, b_j)$ if and only if $i < j$, and hence Theorem 8.1.6 implies that T is unstable. □

Let T be a τ-theory and let $\phi(x_1, \dots, x_n; y_1, \dots, y_m)$ be a first-order τ-formula. If $\underline{M}$ is a model of T, then $A \subseteq M^n$ is *shattered* by ϕ if there exists a family $\{b_I \in M^m \mid I \subseteq A\}$ such that

$$\underline{M} \models \phi(a, b_I) \text{ if and only if } a \in I.$$

By compactness, A is shattered by ϕ if and only if every finite subset of A is shattered by ϕ.

Definition 8.4.3. Let T be a theory. A formula $\phi(x_1, \dots, x_n; \bar{y})$ has the *independence property (IP)* in T if T has a model $\underline{A}$ with an infinite subset of A^n that is shattered by ϕ. Otherwise, we say that ϕ is *dependent* in T. A theory T has the IP if there exists a formula that has the IP in T; otherwise, we say that T has the NIP.

Example 56. The theory of $(\mathbb{N}; \cdot)$ is unstable, witnessed by the formula $\phi(x, y)$ given by

$$\exists k.\, y \cdot k = x$$

(i.e., the formula expresses that x divides y). Then $\phi(x, y)$ shatters the set of prime numbers: for any subset I of prime numbers define $b_I := \prod I$. Then we have $\underline{M} \models \phi(a, b_I)$ if and only if $a \in I$. △

Remark 8.4.4. To prove that ϕ has the IP, the by compactness of first-order logic it suffices to find arbitrarily large finite subsets of A^n that are shattered by ϕ.

Example 57. A typical example of a theory that has the IP is the theory T of the Rado graph $(\mathbb{V}; E)$ (Example 26). Let S be any finite subset of $\mathbb{V}$. Then for any $U \subseteq S$ there exists a vertex $v \in \mathbb{V}$ such that for every $x \in A$ we have $E(x, v) \Leftrightarrow x \in U$. Therefore, the formula $E(x, y)$ shatters S, and Remark 8.4.4 implies that T has the IP. △

Example 58. If T is the theory of an infinite Boolean algebra $(A; 0, 1, \neg, \vee, \wedge)$, then the formula $\phi(x; y) := (x \wedge y = x)$ is independent in T: indeed, let A be a set such that $x \wedge b = 0$ for all distinct $a, b \in A$. For any $I \subseteq N$, set b_I to be $\bigvee_{i \in I} b_i$; then $T \models \phi(x; y)$ if and only if $x \wedge \bigvee_{i \in I} b_i = \bigvee_{i \in I} (x \wedge b_i) = x$ if and only if $x \in I$. △

Proposition 8.4.5. *If a theory T has the IP, then it is unstable.*

Proof. Let $\underline{M} \models T$ and let $A \subseteq M^n$ be infinite and shattered by the formula $\phi(x_1, \dots, x_n, y_1, \dots, y_m)$, witnessed by the family $\{b_I \in M^m \mid I \subseteq A\}$. Let $a_1, a_2, \dots$ be distinct elements of A. Set $c_i := b_{\{a_1, \dots, a_{i-1}\}}$. Then note that $\underline{M} \models \phi(a_i, c_j)$ if and only if $i < j$, so ϕ has the order property. Theorem 8.1.5 then implies that ϕ is unstable in T. □

PROPOSITION 8.4.6. *A formula $\phi(x_1, \ldots, x_n, y_1, \ldots, y_m)$ is NIP in T if and only if for any model $\underline{A} \models T$, every indiscernible sequence $(b_i)_{i \in I}$ of elements of A^m, an every $a \in A^n$ there are at most n indices $i_0 < \cdots < i_{n-1}$ such that for all $i < n-1$*

$$\underline{A} \models \phi(a, b_i) \Leftrightarrow \neg\phi(a, b_{i+1}).$$

PROOF. If ϕ has the IP, then by Lemma 8.2.3 we can choose an indiscernible sequence $(b_i)_{i\in\omega}$ of elements in A^m such that $\{b_i \mid i \in \omega\}$ is shattered by ϕ. Let $s \subseteq \omega$ be the set of even integers. Then $\phi(\bar{x}, b_s)$ satisfies the statement of the lemma: TODO. $\square$

THEOREM 8.4.7. *A theory T is unstable if and only if it has the IP or the SOP.*

PROOF. We have already proved the backwards direction in Proposition 8.4.2 and Proposition 8.4.5.

By Theorem 8.1.6 there exists a formula $\phi(\bar{x}, \bar{y})$, a model $\underline{M} \models T$, and a sequence $(\bar{c}_i)_{i\in\omega}$ such that $\underline{M} \models \phi(\bar{c}_i, \bar{c}_j) \Leftrightarrow i < j$. By Lemma 8.2.3 we may assume that there exists a order indiscernible sequence $(\bar{b}_i)_{i\in\mathbb{Q}}$ such that TODO.

Since T is NIP, Proposition TODO

$\square$

CHAPTER 9

NIP Theories

UNDER CONSTRUCTION.

9.1. VC Dimension

We have already mentioned in Remark 8.4.4 that if $\phi(x_1, \ldots, x_n; \bar{y})$ is dependent in T, then by compactness there is some integer k such that no subset A of M^n of size k is shattered by ϕ. The maximal k for which there is some A of size k that is shattered by ϕ is called the *VC-dimension* of ϕ (called after Vapnik-Chervonenkis [**30**] and discovered in the context of machine learning). If ϕ is independent, we say that the VC-dimension of ϕ is infinite.

Example 59. The formula $x \leq y$ is dependent in $\mathrm{Th}(\mathbb{Q}; \leq)$, and has VC-dimension 1. indeed, $A = \{0\}$ is shattered by ϕ because we can choose $b_A := 0$ and $b_\emptyset := -1$. On the other hand, for $A = \{0, 1\}$ we cannot find a b_1 such that $(\mathbb{Q}; \leq) \models \neg(0 \leq b_1) \wedge 1 \leq b_1$. $\triangle$

Example 60. We have seen in Example 56 that the formula $\phi(x; y) \coloneqq \exists z. x \cdot z = y$ is independent in $T \coloneqq \mathrm{Th}(\mathbb{N}; \cdot)$, and hence has infinite VC dimension. $\triangle$

Example 61. In the countable random graph, every finite subset of vertices is scattered by the formula $E(x, y)$, so this formulas has infinite VC dimension. $\triangle$

If $\phi(\bar{x}; \bar{y})$ is a formula, we write $\phi^-(\bar{y}; \bar{x})$ for $\phi(\bar{x}; \bar{y})$ (so it is the same formula, but we change the role of the variables and parameters).

Lemma 9.1.1. *The formula $\phi(\bar{x}; \bar{y})$ is NIP in T if and only if the formula $\phi^-(\bar{y}; \bar{x})$ is NIP in T.*

Proof. Let $\underline{M}$ be an uncountable model of T. If $\phi(\bar{x}; \bar{y})$ is independent then by compactness, there is an uncountable set $A = \{a_J \mid J \subseteq \omega\}$ which is shattered by ϕ, as witnessed by tuples b_I for $I \subseteq A$. For $j \in \omega$, let $I_j := \{X \subseteq \omega \mid j \in X\}$. We will show that $B := \{b_{I_j} \mid j \in \omega\}$ is shattered by $\phi(\bar{y}; \bar{x})$. Let $K \subseteq \omega$. Then we have

$$
\begin{aligned}
T \models \phi^-(b_{I_j}, a_K) & \text{ if and only if } T \models \phi(a_K; b_{I_j}) \\
& \text{ if and only if } K \in I_j \\
& \text{ if and only if } j \in K\,.
\end{aligned}
$$

Therefore $\phi^-(\bar{y}; \bar{x})$ is independent. $\square$

Exercises.

(93) Show that the VC dimension of $\phi(\bar{x}; \bar{y})$ can be different from the VC dimension of $\phi^-(\bar{y}, \bar{x})$.

Lemma 9.1.2. *The formula $\phi(x; y)$ is independent in T if and only if there is an indiscernible sequence $(a_i \mid i \in \omega)$ and a tuple b such that $T \models \phi(a_i; b)$ if and only if i is even.*

Proof. ($\Leftarrow$): ($\Rightarrow$): Assume that $\phi(x;y)$ in independent in T. Let $A = (a_i \mid i \in \omega)$ be a sequence of $|x|$-tuples that is shattered by $\phi(x;y)$. □

Proposition 9.1.3. *The formula $\phi(x;y)$ is NIP in a theory T if and only if for any model $\underline{M}$ of T and. every indiscernible sequence $(a_i \in M \mid i \in \omega)$ and tuple b, there is some end segment S and $\epsilon \in \{0,1\}$ such that $\phi(a_i;b)^\epsilon$ holds for any $i \in S$.*

Lemma 9.1.4. *Let T be a theory. Then a Boolean combination of NIP formulas is NIP.*

Proof. It is clear from the definition that the negation of a NIP formula is NIP. Let $\phi(x;y)$ and $\psi(x;y)$ be two NIP formulas and we want to show that $\theta(x;y) := \phi \wedge \psi$ is NIP. TODO: finish. □

A τ-theory T is NIP if all τ-formulas ϕ are dependent in T. Note that if T is dependent, then also all formulas with parameters are dependent, since if $\phi(x;y,d)$ has IP, the so does $\phi(x;y,z)$.

Proposition 9.1.5. *Let T be a τ-theory. Then T is NIP if and only if all τ-formulas $\phi(x;y)$ with $|y| = 1$ are dependent.*

Proof. For the interesting direction of the statement, assume that all formulas $\phi(x;y)$ with $|y| = 1$ are dependent. □

We give examples of NIP theories.

Example 62. Any theory of a linear order is NIP. TODO. △

Example 63. Any theory of a tree is NIP. TODO. △

Example 64. The theory $\mathrm{Th}(\mathbb{Q}_p)$ of the p-adic numbers is NIP (either in the language of fields, or in the language of valued fields). TODO. △

Theorem 9.1.6. *Every unstable ω-categorical NIP theory interprets $(\mathbb{Q};<)$.*

Proof. TODO. □

9.2. The Lemma of Sauer-Shelah

Let $\mathcal{F} = \{S_1, S_2, \dots\}$ be a family of sets. Then a set T is *shattered by $\mathcal{F}$* if for every subset $T' \subseteq T$ there exists an $S_i \in \mathcal{F}$ such that $T' = S_i \cap T$. The *VC dimension of $\mathcal{F}$* is the largest cardinality of a set shattered by $\mathcal{F}$.

Theorem 9.2.1 (Sauer-Shelah). *Let $\mathcal{F}$ be a set of subsets of $\{1,\dots,n\}$. If the VC dimension of $\mathcal{F}$ is at most d, then $|\mathcal{F}| \in O(n^d)$.*

We deduce this theorem from a slightly more general theorem which has a very elegant short proof.

Lemma 9.2.2 (Pajor). *Let $\mathcal{F}$ be a finite family of sets. Then there are at least $|\mathcal{F}|$ many sets that are shattered by $\mathcal{F}$.*

Proof. By induction on $|\mathcal{F}|$. For the induction base, note that every family containing only one set shatters the empty set.

For the inductive step, assume the lemma is true for all families of size less than $|\mathcal{F}|$ and let $\mathcal{F}$ be a family of two or more sets. Let x be an element that belongs to some but not all of the sets in $\mathcal{F}$. Split $\mathcal{F}$ into two subfamilies: the sets that contain x and the sets that do not contain x. By inductive assumption, these two subfamilies shatter two collections of sets whose sizes add to at least $|\mathcal{F}|$. None of these shattered sets contain x since a set that contains x cannot be shattered by a family in which all sets contain x or all sets do not contain x. Some of the shattered sets may be shattered by both subfamilies. When a set S is shattered by only one of the two subfamilies then it contributes both to the number of shattered sets of the subfamily and to the number of shattered sets of $\mathcal{F}$. When a set S is shattered by both subfamilies, then both S and $S \cup \{x\}$ are shattered by $\mathcal{F}$, so S contributes twice to the number of shattered sets of the two subfamilies and to the number of shattered sets of $\mathcal{F}$. Therefore, the number of shattered sets of $\mathcal{F}$ is at least equal to the number shattered by the two subfamilies of $\mathcal{F}$, which is at least $\mathcal{F}$. $\square$

Proof of Theorem 9.2.1. Only $\sum_{i=0}^{d-1} \binom{n}{i}$ subsets of $\{1, \ldots, n\} \in O(n^d)$ have cardinality less than d, so if $|\mathcal{F}| \notin O(n^d)$ then by Lemma 9.2.2 there must be a set of cardinality at least k that is shattered, which implies that the VC dimension is larger than d. $\square$

Corollary 9.2.3. *Let T be ω-categorical. Then T is NIP if and only if (orbit growth description).*

CHAPTER 10

Geometric Model Theory

One goal of geometric model theory (or geometric stability theory) is the classification of structures in terms of dimension-like quantities that can be axiomatised using ideas from combinatorial geometry, such as matroids. Key guiding examples include vector spaces and algebraically closed fields. This chapter is not yet finished.

10.1. Pregeometries

Pregeometries and geometries are the key axiomatic notion of geometric model theory; geometries are also known as *matroids* in combinatorics.

DEFINITION 10.1.1. Let X be a set. A *closure operator* on X is a function $C\colon 2^X \to 2^X$ satisfying the following conditions:

(1) C is *extensive* (or *reflexive*), i.e., $U \subseteq C(U)$ for every $U \subseteq X$;
(2) C is *increasing*, i.e., $U \subseteq V \Rightarrow C(U) \subseteq C(V)$ for all $U, V \subseteq X$;
(3) C is *idempotent* (or *transitive*), i.e., $C(C(U)) = C(U)$ for every $U \subseteq X$.

If $C(U) = U$ then $U \subseteq X$ is called *closed*.

We sometimes write $C(x_1, \ldots, x_n)$ instead of $C(\{x_1, \ldots, x_n\})$.

PROPOSITION 10.1.2. *Algebraic closure in a structure $\underline{A}$ (see Definition 6.1.1) is a closure operator on A.*

PROOF. It is obvious that $\mathrm{acl}_{\underline{A}}$ is extensive and increasing. To show that it is idempotent, suppose that $b_1, \ldots, b_k \in \mathrm{acl}_{\underline{A}}(C)$ and that $\underline{A} \models \psi(a, b_1, \ldots, b_k)$ where $\{y \in A \mid \underline{A} \models \psi(y, b_1, \ldots, b_k)\}$ has exactly $m \in \mathbb{N}$ elements. For $i \in \{1, \ldots, k\}$, let $\phi_i(x, \bar{c}_i)$, for a tuple $\bar{c}_i$ of elements from C, be a formula witnessing that $b_i \in \mathrm{acl}_{\underline{A}}(C)$. Let $\phi_{m,\psi}(x_1, \ldots, x_k)$ be a formula that says that $\{y \in A \mid \underline{A} \models \psi(y, x_1, \ldots, x_k)\}$ has at most m elements. Then

$$\exists x_1, \ldots, x_k (\psi(y, x_1, \ldots, x_k) \wedge \phi_{m,\psi}(x_1, \ldots, x_k) \wedge \bigwedge_{i=1}^{k} \phi_i(x_i, \bar{c}_i))$$

is a formula with parameters in C that witnesses that $c \in \mathrm{acl}_{\underline{A}}(C)$, proving idempotency of $\mathrm{acl}_{\underline{A}}$. □

EXAMPLE 65. In algebraically closed fields $\underline{K}$, we have $a \in \mathrm{acl}_{\underline{K}}(A)$ precisely if a is algebraic (in the field-theoretic sense) over the subfield $\underline{F}$ of $\underline{K}$ generated by A (see Appendix B.4.3). This follows from quantifier elimination of ACF (Theorem 7.7.1):

- if a is algebraic in the field-theoretic sense over $\underline{F}$, then there exists a polynomial $p \in F[x]$ such that $p(a) = 0$. Since $\{b \in K \mid \underline{K} \models p(b) = 0\}$ is finite and contains a. This shows that a is algebraic in the model-theoretic sense.
- if $a \in \mathrm{acl}_{\underline{K}}(A)$, there exists a first-order formula $\phi(x)$ such that $\{b \in K \mid \underline{K} \models \phi(b)\}$ is finite and contains a. Since ACF has quantifier elimination, ϕ can be written as a finite disjunction of finite conjunctions of equalities $p(x) = 0$ and inequalities $p(x) \neq 0$. So a must satisfy at least one of the disjuncts. This disjunct cannot consist entirely of inequalities since in this

case infinitely many elements of K would satisfy the disjunct, contradicting our assumptions. Otherwise, if this disjunct contains an equality $p(x) = 0$, then $p(a) = 0$ witnesses field-theoretic algebraicity of a. △

DEFINITION 10.1.3. A *pregeometry* (or *matroid*) on X is a closure operator C on X satisfying the following additional conditions:

(4) C is *finitary*, i.e., if $A \subseteq X$ and $a \in C(A)$, then there is a finite $A' \subseteq A$ such that $a \in C(A')$.
(5) C has the *exchange property*, i.e., if $A \subseteq X$ and $a, b \in X$ then

$$a \in C(A \cup \{b\}) \text{ implies } a \in C(A) \text{ or } b \in C(A \cup \{a\}).$$

A *geometry* (or *simple matroid*) on X is a pregeometry on X such that $C(\emptyset) = \emptyset$ and $C(\{x\}) = \{x\}$ for all $x \in X$.

EXAMPLE 66. Algebraic closure in a structure $\underline{B}$ (see Definition 6.1.1) is clearly always finitary. But it might not have the exchange property. Consider for example the structure $\underline{A} := (A; E, P)$ where

- $A := \mathbb{Q} \times \mathbb{Q}_{\geq 0}$;
- $P := \{(u, v) \in D \mid v = 0\}$;
- $E := \{((u, v), (p, q)) \in D^2 \mid v = 0, q \neq 0, u = p\}$.

It is easy to see that this structure is homogeneous and therefore ω-categorical. Note that $(0, 0)$ is in $\mathrm{acl}_{\underline{A}}(\{(0, 1)\})$, but not in $\mathrm{acl}_{\underline{A}}(\emptyset)$ since $(0, 0)$ lies in the infinite orbit P of $\mathrm{Aut}(\underline{A})$. But $(0, 1)$ is not in $\mathrm{acl}_{\underline{A}}(\{(0, 0)\})$ because there are automorphisms that fix $(0, 0)$ and map $(0, 1)$ to $(0, 2)$, say, showing that $\mathrm{acl}_{\underline{A}}$ fails exchange. △

EXAMPLE 67. Let $\underline{V}$ be a vector space over a field $\underline{F}$. For $A \subseteq V$, let $C(A)$ be the *linear span of* A, i.e., the set of all elements that can be written as linear combinations $u_1 a_1 + \cdots + u_n a_n$ for $a_1, \ldots, a_n \in A$ and $u_1, \ldots, u_n \in F$. (The linear span is the smallest linear subspace of V that contains A.) Then C is an example of a pregeometry. It is not a geometry since $0 \in C(\emptyset)$. △

EXAMPLE 68. Let V be a vector space over a field F. For $A \subseteq V$, let $C(A)$ be the *affine hull of* A, i.e., the set of all elements that can be written as $u_1 a_1 + \cdots + u_n a_n$ where $a_1, \ldots, a_n \in A$ and $u_1 + \cdots + u_n = 1$ $u_1, \ldots, u_n \in F$. The affine hull is the smallest affine subspace of V that contains A. Then C is a geometry. △

EXAMPLE 69. Let V be a vector space over a field F, and let $X = P(V)$ be the projective space over V, i.e, the set of equivalence classes of the equivalence relation $\sim$ defined on $V \setminus \{0\}$ by

$$u \sim v \text{ if and only if } v = \lambda u \text{ for some } \lambda \in F \setminus \{0\}.$$

Recall that a subspace of X is a subset $S \subseteq X$ such that $\bigcup S \cup \{0\}$ is a linear subspace of $\underline{V}$. For $S \subseteq X$, define $C(S)$ to be the smallest subspace of X that contains S. Then C is a geometry. △

EXAMPLE 70. Let $\underline{F}$ be an algebraically closed field. Then $\mathrm{acl}_{\underline{F}}$ (recall that model-theoretic and field-theoretic algebraic closure coincide, see Example 65) is a pregeometry on F. To prove the exchange property, let $A \subseteq F$ be a subset, and let $\underline{K}$ be the subfield of $\underline{F}$ with domain $\mathrm{acl}_{\underline{F}}(A)$. Let $a, b \in F$ be such that $a \in \mathrm{acl}_{\underline{F}}(\{b\} \cup A)$. So there is a non-zero polynomial $p \in \underline{K}[x, y]$ with $p(a, b) = 0$. This implies that $b \in \mathrm{acl}_{\underline{F}}(\{a\} \cup A)$. The operator $\mathrm{acl}_{\underline{F}}$ is not a geometry since $1 \in C(\emptyset)$. △

10.1.1. Bases and dimension. Let C be a pregeometry on X.

DEFINITION 10.1.4. A subset $Y \subseteq X$ is called

- *independent* if $a \notin C(Y \setminus \{a\})$ for all $a \in Y$;
- *generating* if $X = C(Y)$;
- a *basis* if Y is independent and generating.

LEMMA 10.1.5. *Let C be a pregeometry on X with generating set E. Any independent subset of E can be extended to a basis contained in E. In particular, every pregeometry has a basis.*

PROOF. By Zorn's lemma we can choose a maximal independent subset B of E. We claim that $E \subseteq C(B)$ and therefore $C(B) = X$. Suppose otherwise that there is an $x \in E \setminus C(B)$. By the maximality of B the set $B \cup \{x\}$ is not independent, i.e., there is a $b \in B \cup \{x\}$ such that $b \in C(B \cup \{x\} \setminus \{b\})$. This is clearly impossible for $b = x$ since $x \notin C(B)$. If $b \in B$, then $b \notin C(B \setminus \{b\})$ and the exchange property implies that $x \in C(B \cup \{b\} \setminus \{b\}) = C(B)$, a contradiction. □

Since a pregeometry satisfies the Steinitz exchange property all bases are of the same cardinality.

LEMMA 10.1.6. *All bases of a pregeometry C have the same cardinality.*

PROOF. Let I be independent and G a generating subset of X. It suffices to show that $|I| \leq |G|$. Assume first that I is infinite. Then we extend I to a basis B by Lemma 10.1.5. Choose for every $g \in G$ a finite $B_g \subseteq B$ with $g \in C(B_g)$. Since the union of the B_g is a generating set, and B is independent, we have $B = \bigcup_{g \in G} B_g$. This implies that G is infinite, and that $|I| \leq |B| \leq |G|$ since each of the B_g is finite.

Now assume that I is finite. It suffices to show that for every $a \in I \setminus G$ there is some $b \in G \setminus I$ such that $I' = \{b\} \cup I \setminus \{a\}$ is independent. To show this, first observe that G cannot be contained in $C(I \setminus \{a\})$ because $a \in C(G)$, contradicting the independence of I. Choose $b \in G \setminus C(I \setminus \{a\})$. Then the exchange property implies that $\{b\} \cup I \setminus \{a\}$ is independent. □

We can therefore define the *dimension* $\dim(X)$ of X as the size of a basis for X.

EXAMPLE 71. The dimensions of the spaces from Examples 67, 68, 69, and 70 are as follows.

- The linear space of a κ-dimensional vector space has dimension κ.
- The affine space of a κ-dimensional vector space has dimension $\kappa + 1$.
- The projective space of a κ-dimensional vector space has dimension κ.
- The dimension of an algebraically closed field $\underline{F}$ with respect to $\mathrm{acl}_{\underline{F}}$ equals the transcendence degree of $\underline{F}$ over the prime field. △

Exercises.

(94) Let C be a pregeometry over X. Show that the following are equivalent for all $B \subseteq X$:
  - B is a basis for X.
  - B is a maximal independent set.
  - B is a minimal set that generates X.

10.1.2. Restriction and Relativisation. Any subset $Y \subseteq X$ of a pregeometry C on X gives rise to two new pregeometries, the *restriction* and the *relativisation* of C with respect to Y. The *restriction* of C to Y is the pregeometry C^Y on Y defined by

$$C^Y(S) := C(S) \cap Y$$

for all $S \subseteq Y$. It is immediate that $A \subseteq Y$ is independent with respect to C if and only if it is independent with respect to C^Y, and that it is a basis for Y with respect to C if and only if it is a basis for Y with respect to C^Y.

DEFINITION 10.1.7. Let C be a pregeometry on X. The dimension of $Y \subseteq X$ is the dimension of the pregeometry C^Y on Y.

The *relativisation* of C to Y is the pregeometry C_Y on X defined by

$$C_Y(S) := C(S \cup Y)$$

for all $S \subseteq X$.

PROPOSITION 10.1.8. *Let C be a pregeometry on X and $Y \subseteq X$. Let A be a basis of C^Y and B a basis of C_Y. Then $A \cup B$ is a basis of C.*

PROOF. We have $C^Y(A) = Y$ and $C_Y(B) = X$. Therefore, $X = C_Y(B) = C(B \cup Y) = C(B \cup (C(A) \cap Y)) \subseteq C(B \cup A)$ and hence $A \cup B$ is generating. Since B is independent with respect to C_Y, we have $b \notin C_Y(B \setminus \{b\}) = C(A \cup B \setminus \{b\})$ for all $b \in B$. Consider an $a \in A$. We have to show that $a \notin C(A' \cup B)$ where $A' = A \setminus \{a\}$. As $a \notin C(A')$ we let B' be a maximal subset of B with $a \notin C(A' \cup B')$. If $B' \neq B$ this would imply that $a \in C(A' \cup B' \cup \{b\})$ for any $b \in B \setminus B'$ which would in turn imply $b \in C(A \cup B')$, a contradiction. □

10.1.3. The Associated Geometry. There is a standard way of obtaining a geometry from a pregeometry C on X. First, if $C(\emptyset) \neq \emptyset$, then replace X by

$$X' := X \setminus C(\emptyset),$$

equipped with the restriction pregeometry C'. Then define an equivalence relation $\sim$ on X' by defining $u \sim v$ if $C'(x) = C'(y)$, and define a pregeometry C'' on $X'' := X'/{\sim}$ by

$$C''(S) := \{[b]_\sim \mid b \in C'(\bigcup S)\}$$

for all $S \subseteq X''$. Then C'' is indeed a geometry, called the *geometry associated to C*. This generalises the way in which we obtain projective spaces from vector spaces.

10.1.4. Special Types of Pregeometries. A pregeometry C on X is called

- *trivial* if $C(A) = \bigcup_{a \in A} C(\{a\})$.
- *modular* if for any two closed $A, B \subseteq X$

$$\dim(A \cup B) = \dim(A) + \dim(B) - \dim(A \cap B).$$

- *locally modular* if $\dim(A \cup B) = \dim(A) + \dim(B) - \dim(A \cap B)$ holds for all closed sets $A, B \subset X$ with $\dim(A \cap B) > 0$.
- *locally finite* if closures of finite sets are finite.

EXAMPLE 72. Structures $\underline{A}$ with no algebraicity (i.e., $\mathrm{acl}_{\underline{A}}(S) = S$ for all $S \subseteq A$) are trivial. Less trivial, but still trivial in the sense of the definition above is the structure that contains an equivalence relation with infinitely many classes of size $k \in \mathbb{N}$, which has algebraicity for $k > 1$. △

EXAMPLE 73. The linear span in vector spaces is a prototypical example of a modular pregeometry. Affine spaces are not modular. For example, if A and B are parallel lines in V then $\dim(A) + \dim(B) + \dim(A \cap B) = 2 + 2 + 0$ which is strictly bigger than $\dim(A \cup B) = 3$. However, affine spaces are locally modular. The linear span and affine hull in a vector space are locally finite if and only if the underlying field is finite. △

LEMMA 10.1.9. *C is locally modular if and only if for all $p \in X \setminus C(\emptyset)$ the relativised pregeometry $C_{\{p\}}$ is modular.*

PROOF. Let $p \in X \setminus C(\emptyset)$. Let $A, B \subseteq X$ be closed with respect to $C_{\{p\}}$. Then $A \cap B$ contains p, and since $p \notin C(\emptyset)$ the dimension of $A \cap B$ with respect to C is at least one. The statement now follows from Proposition 10.1.8. □

EXAMPLE 74. Suppose that $\underline{K} \models$ ACF has transcendence degree at least five over the prime field. Then $\mathrm{acl}_{\underline{K}}$ is not locally modular. By Lemma 10.1.9 it suffices to show that for some $p \in X \setminus \mathrm{acl}_{\underline{K}}(\emptyset)$ the relativised pregeometry $C := (\mathrm{acl}_{\underline{K}})_{\{p\}}$ is not modular. Choose an independent set $\{p, a, b, c, d\}$, set $x := \frac{a-c}{b-d}$, and let $A := C(a, b)$ and $B := C(x, a - b)$. Then $C(A \cup B) = C(a, b, x)$ has dimension 3, while $\dim(A) = \dim(B) = 2$. To show that C is not modular it is sufficient to prove that $\dim(A \cap B) \neq 1$.

Claim. $A \cap B = C(\emptyset)$. Let $z \in A \cap B$. We claim that there exists an $\alpha \in \mathrm{Aut}(\underline{K})$ that maps (a, b, c, d) to (c, d, a, b) and that fixes B pointwise. First note that any α that maps (a, b, c, d) to (c, d, a, b) satisfies $\alpha(x) = x$ and

$$\alpha(a - bx) = c - dx = \frac{cb - cd - da + dc}{b - d} = \frac{ab - da - ba + bc}{b - d} = a - bx.$$

Hence, α fixes $\{x, a - bx\}$, and can be extended to to an automorphism of $\underline{K}$ that fixes $B = C(x, a - bx)$ pointwise. In particular, $\alpha(z) = z$. Furthermore, $z \in A = C(a, b)$ implies $z = \alpha(z) \in C(c, d)$. Consequently, $z \in C(a, b) \cap C(c, d) = C(\emptyset)$. △

The sets A and B are *independent over C* if for all sets $A' \subseteq A$ and $B' \subseteq B$, if each A' and B' is independent over C, then A' and B' are disjoint and their union is independent over C. Note that this relation is symmetric. The following is then easy to see.

LEMMA 10.1.10. *For a pregeometry C over X the following are equivalent.*

- *C is modular.*
- *Any two closed A and B are independent over $A \cap B$.*
- *For any two closed sets A and B we have $\dim_B(A) = \dim_{A\cap B}(A)$.*

Exercises.

(95) Show that every pregeometry is *submodular*, i.e.,

$$\dim(A \cup B) + \dim(A \cap B) \leq \dim(A) + \dim(B).$$

(96) The set of all closed subsets of a closure operator forms a lattice, where the infimum $\wedge$ is intersection and the supremum $\vee$ of X and Y is $C(X \cup Y)$. Show that a pregeometry is modular if and only if the lattice of closed sets is modular, i.e., for all closed A, B, C

$$A \subseteq C \Rightarrow A \vee (B \wedge C) = (A \vee B) \wedge C.$$

(97) Show that every trivial pregeometry is modular.
(98) Show that a pregeometry C on X is modular if and only if for all $c \in C(A \cup B)$ there are $a \in C(A)$ and $b \in C(B)$ such that $c \in C(\{a, b\})$;
(99) Prove that a geometry C on X is modular if and only if whenever $a, b \in X$ and $A \subseteq X$ such that $\dim(\{a, b\}) = 2$ and $\dim_A(\{a, b\}) \leq 1$ then

$$C(\{a, b\} \cap C(A)) \setminus C(\emptyset) \neq \emptyset.$$

10.2. Minimal and Strongly Minimal Sets

Definition 10.2.1 (Minimal sets). Let $\underline{B}$ be a structure and let $D \subseteq B$ be definable in $\underline{B}_B$. Then

- D is *minimal in* $\underline{B}$ if the only subsets of D that are definable in $\underline{B}_B$ are finite or cofinite in D.
- If $\phi(x)$ is a formula that defines a minimal set D over $\underline{B}_B$, then we also say that ϕ is *minimal in* $\underline{B}$.
- The structure $\underline{B}$ is called *minimal* if if B is minimal in $\underline{B}$.

Exercises.

(100) Show that $(\mathbb{Q}; <)$ and the random graph are not minimal.

Example 75. Let E be an equivalence relation on a countable set S with infinitely many two-element classes. Then $(S; E)$ is minimal. $\triangle$

In fact, the structure from Example 75 satisfies a strengthened from of minimality, defined next.

Definition 10.2.2 (Strongly minimal sets). Let $\underline{B}$ be a structure. Then $D \subseteq B$ is *strongly minimal in* $\underline{B}$ if it is minimal in every elementary extension of $\underline{B}$. A theory T is *strongly minimal* if for every model $\underline{B}$ of T the underlying set B is strongly minimal. A structure $\underline{B}$ is strongly minimal if its first-order theory is strongly minimal.

The claims for the following examples follow easily from quantifier elimination.

Example 76. Let V be the countably infinite vector space over the two-element field $\mathbb{F}_2$. Then $(V; +)$ is strongly minimal. $\triangle$

Example 77. Let $\underline{K}$ be an algebraically closed field. Then $\underline{K}$ is strongly minimal. $\triangle$

Example 78. The structure $(\mathbb{N}; <)$ is minimal, but not strongly minimal since in every elementary extension $\underline{B}$ there is an element $c \in B \setminus \mathbb{N}$ such that $(x < c)^{\underline{B}}$ and $(x > c)^{\underline{B}}$ are both infinite. $\triangle$

Theorem 10.2.3 (Baldwin, Lachlan). *Let $\underline{B}$ be a structure and let $X \subseteq B$ be minimal in $\underline{B}$. Then the restriction of* $\mathrm{acl}_{\underline{B}}$ *to X is a pregeometry.*

Proof. Let $A \subseteq X$ and let $c \in \mathrm{acl}(A \cup \{b\}) \setminus \mathrm{acl}(A)$ for some $b, c \in X$; we want to show that $b \in \mathrm{acl}(A \cup \{c\})$. By assumption, there is a formula $\phi(x, y)$ with parameters from A such that $\underline{B} \models \phi(c, b)$ and $|\{x \in X \mid \underline{B} \models \phi(x, b)\}| = n$ for some $n \in \mathbb{N}$. Let $\psi(y)$ be a formula with parameters from A that says that $|\{x \in X \mid \underline{B} \models \phi(x, y)\}| = n$. Then $\underline{B} \models \psi(b)$. If $S := \{y \in X \mid \underline{B} \models \phi(c, y) \wedge \psi(y)\}$ is finite, then $\phi(c, y) \wedge \psi(y)$ would witness that $b \in \mathrm{acl}(A \cup \{c\})$ and we are done. Otherwise, S is cofinite by the minimality of $\underline{B}$, so that $|X \setminus S| = m$ for some $m \in \mathbb{N}$. Let $\chi(x)$ be a formula with parameters from A stating that

$$|\{y \in X \mid \underline{B} \models \neg(\phi(x, y) \wedge \psi(y))\}| = m.$$

If $\chi(x)$ defines a finite subset of X, then $c \in \mathrm{acl}(A)$, a contradiction. Hence, $\chi(x)$ defines a cofinite subset of X. We may therefore choose $n + 1$ distinct elements $a_1, \dots, a_{n+1} \in X$ such that $\underline{B} \models \chi(a_i)$. This means that

$$B_i := \{u \in X \mid \underline{B} \models \phi(a_i, u) \wedge \psi(u)\}$$

is cofinite for $i \in \{1, \dots, n + 1\}$. Hence $\bigcap_{i \in \{1,\dots,n+1\}} B_i$ contains an element b'. We have at least $n+1$ elements $x \in X$ such that $\phi(x, b')$ is satisfied, namely $a_1, \dots, a_{n+1}$. But this contradicts the fact that $\psi(b')$. $\square$

In the following, the notion of dimension in the context of minimal subsets is always with respect to the pregeometry from algebraic closure (also see Definition 10.1.7).

LEMMA 10.2.4. *Let* $\underline{A}$ *and* $\underline{B}$ *be minimal structures with the same first-order theory, let* $A' \subseteq A$ *and* $B' \subseteq B$*. If* $\dim(A') = \dim(B')$*, then any bijection between a basis for* A' *and a basis for* B' *can be extended to an isomorphism between the substructure of* $\underline{A}$ *on* $\mathrm{acl}_{\underline{A}}(A')$ *and the substructure of* $\underline{B}$ *on* $\mathrm{acl}_{\underline{B}}(B')$*.*

PROOF. Let f be a bijection between a basis for A' and a basis for B'. We first show that f preserves all first-order formulas $\phi(x_1, \ldots, x_n)$ by induction on n. If $n = 0$ then the statement holds by the assumption that $\underline{A}$ and $\underline{B}$ have the same first-order theory. Let $a_1, \ldots, a_n$ be elements from the basis for A' such that $\phi(a_1, \ldots, a_n)$. Since $a_1, \ldots, a_n$ are independent, we have that $a_n \notin \mathrm{acl}_{\underline{A}}(\{a_1, \ldots, a_{n-1}\}$. Since $\underline{A}$ is minimal, the formula $\phi(a_1, \ldots, a_{n-1}, x_n)$ is satisfied by cofinitely many x_n in A. In other words, there exists an $\ell \in \mathbb{N}$ such that $\underline{A} \models \exists x^{=\ell} x_n. \neg\phi(a_1, \ldots, a_{n-1}, x_n)$ (see Exercise 8). By the inductive assumption, $\underline{B} \models \exists x^{=\ell} x_n. \neg\phi(f(a_1), \ldots, f(a_{n-1}), x_n)$. Since $f(a_1), \ldots, f(a_n)$ are independent, there is no formula satisfied by $f(a_{n+1})$ and finitely other elements from the basis of B'. We conclude that $\underline{B} \models \phi(f(a_1), \ldots, f(a_n))$.

Next, we show how to extend f to an isomorphism between the substructure of $\underline{A}$ on $A'' := \mathrm{acl}_{\underline{A}}(A')$ and the substructure of $\underline{B}$ on $B'' := \mathrm{acl}_{\underline{B}}(B')$. Suppose inductively that we have already constructed a map $g \colon E \to B''$, for $E \subseteq A''$ with $A' \subseteq E$, that preserves all first-order formulas, and let $a \in A'' \setminus A'$. Since $a \in \mathrm{acl}_{\underline{A}}(A')$ there is a formula $\phi(x, y_1, \ldots, y_n)$ and $a_1, \ldots, a_n \in E$ such that $\underline{A} \models \exists^{=\ell} x.\phi(x, a_1, \ldots, a_n)$ for some $\ell \in \mathbb{N}$. Choose ϕ and $a_1, \ldots, a_n \in E$ such that ℓ is as small as possible. By assumption, $\underline{N} \models \exists^{=\ell} x.\phi(x, g(a_1), \ldots, g(a_n))$. In particular there exists $b \in B''$ such that $\underline{N} \models \phi(b, g(a_1), \ldots, g(a_n))$. Define the extension g' of g by setting $g'(a) := b$.

To show that g' preserves all first-order formulas, let $\psi(x, y_1, \ldots, y_m)$ be a formula and $d_1, \ldots, d_m \in E$. Suppose that $\underline{A} \models \phi(a, d_1, \ldots, d_m)$. By the minimality of ℓ, we have that $\underline{A} \models \forall x(\phi(x, a_1, \ldots, a_n) \Rightarrow \psi(x, d_1, \ldots, d_m))$. Since $\underline{B} \models \phi(b, g(a_1), \ldots, g(a_n))$, this implies that $\underline{B} \models \psi(b, g(a_1), \ldots, g(a_n))$.

By transfinite induction, we may therefore obtain an embedding $g' \colon A'' \to B''$ which extends f. It remains to show that g' is surjective. Let $c \in B'' = \mathrm{acl}_{\underline{B}}(B')$. Then there is a formula $\phi(x, c_1, \ldots, c_n)$ for some elements $c_1, \ldots, c_n$ from the basis of B' witnessing that $c \in \mathrm{acl}_{\underline{B}}(B')$. Let $\ell \in \mathbb{N}$ be such that $\underline{B} \models \exists^{=\ell} x.\phi(x, c_1, \ldots, c_n)$. For each c_i let $a_i :- f^{-1}(c_i)$. Since f preserves all first-order formulas we obtain that $\underline{A} \models \exists^{=\ell} x.\phi(x, a_1, \ldots, a_n)$. The ℓ distinct elements of B'' that satisfy $\phi(x, c_1, \ldots, c_n)$ are mapped by g' to the ℓ distinct elements of A'' that satisfy $\phi(x, c_1, \ldots, c_n)$. In particular, there exists $a \in A''$ such that $g'(a) = c$. This proves that g' is an isomorphism. □

COROLLARY 10.2.5. *Let* $\underline{A}$ *be a countable minimal structure without algebraicity. Then* $\underline{A}$ *is preserved by all permutations, i.e.,* $\underline{A}$ *is interdefinable with* $(\mathbb{N}; =)$*.*

Exercises.

(101) Show that an ω-categorical structure is strongly minimal if and only if it is minimal. How about the same statement for $\aleph_1$-categorical structures?

(102) In which of the following structures $\underline{B}$ is the algebraic closure operator a pregeometry on B?
- An equivalence relation on a countably infinite set with infinitely many classes of size two
- The random graph
- $(\mathbb{Z}; +)$

10.3. Uncountably Categorical Structures

We have already seen that ACF_p, for $p = 0$ or p prime, is κ-categorical for all uncountable κ. Another example are equivalence relations on a countably infinite set with infinitely many classes of size two. The following example presents some structures that are ω-categorical, but not κ-categorical for uncountable κ.

EXAMPLE 79. Let $(D; A)$ be a structure where D is countably infinite and A is such that both A and $D \setminus A$ are infinite. Then $T := \mathrm{Th}(D; A)$ has models of cardinality $\kappa > \omega$ where $|A| = \kappa$ and $|D \setminus A| = \omega$, and models where $|A| = \omega$ and $|D \setminus A| = \kappa$, and these are not isomorphic, so T is not κ-categorical.

Similarly, one can show that $\mathrm{Th}(D; E)$, where E is an equivalence relation with two infinite classes, is not κ-categorical for $\kappa > \omega$. △

LEMMA 10.3.1. *Let τ be a countable signature and let T be a complete and strongly minimal τ-theory. Then T is κ-categorical for every uncountably cardinal κ.*

PROOF. By Lemma 10.2.4, it is enough to show that any two models of T of cardinality κ have the same dimension. Let $\underline{A}$ be a model of T of cardinality $\kappa > \omega$ and let $B \subseteq A$ be a basis of $\underline{A}$. Since τ is countable, there are only countably many τ-formulas. In particular, there are only countably many τ-formulas over parameters from B that are satisfied by finitely many elements, and it follows that $|\operatorname{acl}_{\underline{A}}(B)| \leq |B|$. Since $\operatorname{acl}_{\underline{A}}(B) = A$ we conclude that $|B| = |A|$. □

We mention a few highlights about κ-categorical structures. We have already seen in Theorem 8.3.2 that κ-categoricity implies ω-stability.

THEOREM 10.3.2 (Morley). *Let κ be an uncountable cardinal. Then T is κ-categorical if and only if it is $\aleph_1$-categorical.*

An important step in the proof of Theorem 10.3.2 is the following partial converse of Lemma 10.3.1.

LEMMA 10.3.3. *Every uncountably categorical structure contains a strongly minimal set.*

DEFINITION 10.3.4. Let T be a τ-theory. Two structures $\underline{A}, \underline{B} \models T$ are called a *Vaughtian pair* if there exists a $(\tau \cup A)$-formula $\phi(x)$ such that

- $\underline{B}$ is a proper elementary extension of $\underline{A}$;
- $\phi^{\underline{A}}$ is infinite;
- $\phi^{\underline{A}} = \phi^{\underline{B}}$.

EXAMPLE 80. The theory of the structure $(D; E)$ from Example 79 has the Vaughtian pair $(D; A), (D \cup E; A)$ where $|E| = \aleph_1$. △

THEOREM 10.3.5 (Baldwin-Lachlan). *Let κ be an uncountable cardinal. A countable theory T is κ-categorical if and only if T is ω-stable and has no Vaughtian pairs.*

Exercises.

(103) (*) Show that the theory of the random graph is not $\aleph_1$-categorical.

APPENDIX A

Set Theory

In this appendix chapter we freely follow Appendix A of the book of Tent and Ziegler [**29**]. See [**20**] for a more detailed introduction to set theory.

A.1. Sets and Classes

We often speak of the *class of all τ-structures* or the *class of all ordinals*, since we know that these things cannot be sets (as we will see in the next section, if the class of all ordinals were a set, we could derive a contradiction). Using first-order logic, we therefore want to give an axiomatic treatment of set theory that allows for the distinction between sets and (proper) classes. For this, we may work in Bernays-Gödel set theory (BG) which is formulated in a two-sorted language. One type of objects are *sets* and the other type of objects are *classes*; sets can be elements of sets, and sets can be elements of classes, but classes can't be elements of sets or classes.

Here are the axioms of BG:

(1) (a) *Extensionality.* Sets containing the same elements are equal:

$$\forall x, y \in \text{Sets}\,\big(\forall z(z \in x \Leftrightarrow z \in y) \Rightarrow x = y\big)$$

(b) *Empty set.* There exists an empty set (denoted by $\emptyset$):

$$\exists x \in \text{Sets}\,\forall y\,(\neg y \in x)$$

(c) *Pairing.* For all sets a and b there is a set (denoted by $\{a, b\}$) which has exactly the elements a and b:

$$\forall a, b \in \text{Sets}\;\; \exists c \in \text{Sets}\;\; \forall x\,(x \in c \Leftrightarrow x = a \vee x = b)$$

(d) *Union.* For every set a there is a set (denoted by $\bigcup a$) that contains precisely the elements of the elements of a:

$$\forall a \in \text{Sets}\,\exists b \in \text{Sets}\,\forall x\,(x \in b \Leftrightarrow \exists y \in a.\, x \in y)$$

(e) *Power set.* For every set a there is a set (denoted by $\mathcal{P}(a)$) that consists of all subsets of a:

$$\forall a \in \text{Sets}\,\exists b \in \text{Sets}\,\forall x\,(x \in b \Leftrightarrow x \subseteq a)$$

(f) *Infinity.* There is an infinite set. One way to express this is to assert the existence of a set which contains the empty set and is closed under the successor operation $x \mapsto x \cup \{x\}$:

$$\exists a \in \text{Sets}\,\big(\emptyset \in a \wedge \forall x\,(x \in a \Rightarrow x \cup \{x\} \in a\big)$$

(2) (a) *Class extensionality*: Classes containing the same elements are equal.

(b) *Comprehension*: If $\phi(x, y_1, \ldots, y_m, z_1, \ldots, z_n)$ is a first-order formula in which only set-variables are quantified, and if $b_1, \ldots, b_m$ are sets, and $c_1, \ldots, c_n$ are classes, then there exists a class, denoted by

$$\{x \in \text{Sets} \mid \phi(x, b_1, \ldots, b_m, c_1, \ldots, c_n\}$$

containing precisely those sets x that satisfy $\phi(x, b_1, \dots, b_m, c_1, \dots, c_n)$. This is in fact an infinite family of axioms (for every first-order formula ϕ we have one axiom).

(c) *Replacement*: If a class c is a function, i.e., if for every set a there is a unique set $b =: c(a)$ such that $(a, b) := \{\{a\}, \{a, b\}\}$ belongs to c, then for every set d the image $\{c(z) \mid z \in d\}$ is a set.

(3) *Regularity* (or *Axiom of Foundation*):

$$\forall x \in \text{Sets}\ \big(x \neq \emptyset \Rightarrow \exists y \in x\, (y \cap x = \emptyset)\big)$$

Consequently: no set can be an element of itself. Moreover, there is no infinite sequence (a_n) such that a_{i+1} is an element of a_i for all i. Most parts of mathematics can be practised without this axiom, but assuming this axiom simplifies some proofs of fundamental properties of ordinals.

(4) For BGC we add *Global Choice*: There is a function c such that $c(a) \in a$ for every nonempty set a.

We mention that BG is a *conservative extension* of ZF: any set-theoretical statement provable in BG (BGC) is also provable in ZF (ZFC). (The substructure of a model of BG induced by sets is a model of ZF. Conversely, a model M of ZF becomes a model of BG by taking the first-order definable subsets of M as classes.)

A.2. Zorn's Lemma

The axiom of choice implies Zorn's lemma, which will be used several times in this text.

Theorem A.2.1 (Zorn's Lemma). *Let $(P; \leq)$ be a partially ordered set with the property that every chain in P has an upper bound in P. Then P contains at least one maximal element.*

A.3. Ordinals

A *well-ordering* of a set A is a linear order $<$ on A such that any non-empty subset of A contains a smallest element with respect to $<$. The usual ordering of the set $\mathbb{N}$ is an example of a well-ordering, and the usual ordering of $\mathbb{Z}$ and the usual ordering of the non-negative rational numbers are non-examples. For $a \in A$, the *predecessors* of a are the elements $b \in A$ with $b < a$.

Definition A.3.1. An *ordinal* is a well-ordered set in which every element equals its set of predecessors.

The well-ordering of an ordinal α is given by the relation $\in$: if $\beta, \gamma \in \alpha$, then

$$\begin{aligned} \beta < \alpha &\Leftrightarrow \beta \text{ is a predecessor of } \alpha \\ &\Leftrightarrow \beta \in \alpha \end{aligned}$$

Hence, an ordinal is uniquely given by the set of its elements. Suppose that α is an ordinal, and $\gamma \in \beta \in \alpha$. Then γ is a predecessor of β, and hence must be in α. That is, ordinals α are *transitive*: if $\beta \in \alpha$ then $\beta \subseteq \alpha$.

An element β of an ordinal α is again an ordinal number: Since $\beta \subseteq \alpha$, we obtain a well-ordering of β by restricting the well-ordering of α to β. Moreover, if $\gamma, \gamma' \in \beta$, then $\gamma < \gamma'$ iff $\gamma \in \gamma'$ so β is indeed an ordinal.

Proposition A.3.2. *Every structure $(A; <)$ where A is well-ordered by $<$ is isomorphic to $(\alpha; \in)$ where α is an ordinal.*

Proof. Define $f\colon A \to \alpha$ inductively by $f(y) := \{f(z) \mid z < y\}$. The image of f is an ordinal α such that $(\alpha; \in)$ is isomorphic to $(A; <)$ via f. (Note that f is the only isomorphism between $(A; <)$ and $(\alpha; \in)$.) □

We denote the class of all ordinals by On.

Proposition A.3.3. On *is a* proper *class, i.e., a class which is not a set, and it is well-ordered by* $\in$.

Proof. Let α and β be different ordinals. We have to show that either $\alpha \in \beta$ or $\beta \in \alpha$. If not, then $x = \alpha \cap \beta$ would be a proper initial segment of α and β, and therefore itself an element of α and β, so $x \in x$, a contradiction.

The class of all ordinals is not a set (the Burali-Forti Paradox): if it were a set, then it were an ordinal itself, and thus a member of itself, contradicting the Axiom of Foundation. □

So every ordinal α is a set having as elements precisely the smaller ordinals:

$$\alpha = \{\beta \in \mathrm{On} \mid \beta < \alpha\} .$$

In the following, for ordinals α, β we write $\alpha < \beta$ rather than $\alpha \in \beta$.

For any ordinal α its *successor* is defined as $\alpha + 1 := \alpha \cup \{\alpha\}$: it is the smallest ordinal greater than α. Starting from the smallest ordinal $0 = \emptyset$, its successor is $1 = \{0\}$, then $2 = \{0, 1\}$, and so on, yielding the natural numbers $\mathbb{N}$. When we view $\mathbb{N}$ as an ordinal, we denote it by ω. The next ordinal is $\omega + 1 = \{0, 1, \ldots, \omega\}$, etc. By definition, a successor ordinal β contains a maximal element α (so $\beta = \alpha + 1$). Ordinals greater than 0 which are not successor ordinals are called *limit ordinals.* These are precisely the ordinals γ that can be written as $\bigcup_{\beta<\gamma} \beta$. Any ordinal can be written uniquely as

$$\lambda \underbrace{+1 + \cdots + 1}_{n \text{ times}}$$

where λ is a limit ordinal.

Theorem A.3.4 (Well-ordering theorem). *Every set has a well-ordering.*

Proof. Let A be a set. Fix a set B which does not belong to A and define a function f from the class of all ordinals to $A \cup \{B\}$ as follows:

- if $A \setminus \{f(\beta) \mid \beta < \alpha\} \neq \emptyset$ then set $f(\alpha)$ to be an element from this set (here we use the Axiom of Choice).
- Otherwise, $f(\alpha) := B$.

Then $\gamma := \{\alpha \mid f(\alpha) \neq B\}$ is an ordinal and f defines a bijection between γ and A. □

In fact, the well-ordering theorem is equivalent to the Axiom of Choice. Note that the ordinal γ in the construction of the well-ordering of A is not unique, unless A is finite.

A.4. Cardinals

Two sets A and B *have the same cardinality* ($|A| = |B|$) if there exists a bijection between them. By the well-ordering theorem, every set has the same cardinality as some ordinal. We call the smallest such ordinal the *cardinality* $|A|$ of A. Ordinals occurring in this way are called *cardinals.* An ordinal α is a cardinal if and only if all smaller ordinals do not have the same cardinality.

Notes.

- All natural numbers and ω are cardinals.

- $\omega + 1$ is the smallest ordinal that is not a cardinal.
- The cardinality of a finite set is a natural number.
- A set of cardinality ω is called *countably infinite*.

Sums, products, and powers of cardinals are defined as the cardinality of disjoint unions, Cartesian powers, and sets of functions:

$$\begin{aligned} |x| + |y| &:= |x \uplus y| && \text{where } x \cap y = \emptyset \\ |x| \cdot |y| &:= |x \times y| \\ |x|^{|y|} &:= |x^y| \end{aligned}$$

and likewise for infinite sums and products:

$$\begin{aligned} \sum_{x \in I} |x| &:= \Big| \biguplus_{x \in I} x \Big| \\ \prod_{x \in I} |x| &:= \Big| \prod_{x \in I} x \Big|. \end{aligned}$$

Note that

$$(\kappa^\lambda)^\mu = \kappa^{\lambda \cdot \mu}.$$

Cantor's well-known diagonalization argument shows that

$$2^\kappa > \kappa\,.$$

In particular, there is no largest cardinal. Cantor's result also follows from König's theorem below for $\kappa_i := 1$ and $\lambda_i := 2$ for all $i \in I := \omega$.

THEOREM A.4.1 (König's theorem). *Let $(\kappa_i)_{i\in I}$ and $(\lambda_i)_{i\in I}$ be sequences of cardinals. If $\kappa_i < \lambda_i$ for all $i \in I$, then*

$$\sum_{i \in I} \kappa_i < \prod_{i \in I} \lambda_i\,.$$

PROOF. We first show that $\sum_{i\in I} \kappa_i \leq \prod_{i \in I} \lambda_i$. Choose pairwise disjoint sets $(A_i)_{i\in I}$ and $(B_i)_{i \in I}$ such that $|A_i| = \kappa_i$, $|B_i| = \lambda_i$, and $A_i \subset B_i$ for all $i \in I$. We will construct an injection $f\colon \bigcup_{i\in I} A_i \to \prod_{i\in I} B_i$. Choose $d_i \in B_i \setminus A_i$ for each $i \in I$ (here we use the Axiom of Choice). For $x \in A := \bigcup_{i \in I} A_i$, define

$$f(x) := (a_i)_{i\in I} \text{ where } a_i := \begin{cases} x & \text{if } x \in A_i \\ d_i & \text{otherwise.} \end{cases}$$

To show the injectivity of f, let $x, y \in A$ be distinct. Let $i \in I$ be such that $x \in A_i$. If $y \in A_i$ then $f(x)_i = x \neq y = f(y)_i$. If $y \notin A_i$ then $f(x)_i = x \neq d_i = f(y)_i$ since $x \in A_i$ but $d_i \in B_i \setminus A_i$. So in both cases, $f(x) \neq f(y)$.

Suppose for contradiction that $\sum_{i\in I} \kappa_i = \prod_{i\in I} \lambda_i$. Then we can find sets $(X_i)_{i\in I}$ with $|X_i| = \kappa_i$ such that

$$B := \prod_{i\in I} B_i = \bigcup_{i \in I} X_i\,.$$

For each $i \in I$, define

$$Y_i := \{a_i \mid a \in X_i\}\,.$$

For every $i \in I$ there exists $b_i \in B_i \setminus Y_i$ because $|Y_i| \leq |X_i| = \kappa_i < \lambda_i = |B_i|$. Now define

$$b := (b_i)_{i\in I} \in \prod_{i\in I} B_i.$$

Let $j \in I$. Then $b_j \notin Y_j$ by the choice of b_j, and hence $b \notin X_j$ by the definition of Y_j. This shows that $b \notin \bigcup_{i\in I} X_i$, a contradiction. □

We write κ^+ for the smallest cardinal greater than κ, the *successor cardinal* of κ. Positive cardinals which are not successor cardinals are called *limit cardinals.* There is an isomorphism between the class of ordinals and the class of all infinite cardinals, which is denoted by

$$\alpha \mapsto \aleph_\alpha$$

and can be defined inductively by

$$\aleph_\alpha := \begin{cases} \omega & \text{if } \alpha = 0 \\ \aleph_\beta^+ & \text{if } \alpha = \beta + 1 \\ \bigcup_{\beta<\alpha} \aleph_\beta & \text{if } \alpha \text{ is a limit ordinal.} \end{cases}$$

Note that if $(\kappa_i)_{i\in I}$ is a family of cardinals, then $\kappa := \bigcup_{i\in I} \kappa_i$ is again a cardinal:

- if there is an $i \in I$ such that $\kappa_j \leq \kappa_i$ for all $j \in I$, then $\bigcup_{i\in I} \kappa_i = \kappa_i$ and the statement is true;
- otherwise, for every $i \in I$ there is a $j \in I$ with $\kappa_i < \kappa_j$. For each ordinal α with $\alpha < \kappa$ we have that $\alpha \in \kappa$ and hence $\alpha \in \kappa_i$ for some $i \in I$. By the above, there is a $j \in I$ such that $|\alpha| \leq \kappa_i < \kappa_j \leq |\kappa|$. Thus, every ordinal smaller than κ has smaller cardinality than κ, and κ is a cardinal.

THEOREM A.4.2. *Let κ be an infinite cardinal. Then*

(1) $\kappa \cdot \kappa = \kappa$.
(2) $\kappa + \lambda = \max(\kappa, \lambda)$.
(3) $\kappa^\kappa = 2^\kappa$.

PROOF. For ordinals $\alpha, \beta, \alpha', \beta'$, define $(\alpha, \beta) < (\alpha', \beta')$ iff

$$(\max(\alpha, \beta), \alpha, \beta) <_{\text{lex}} (\max(\alpha', \beta'), \alpha', \beta')$$

where lex is the lexicographical ordering on triples of ordinals. Since this is a well-ordering, there is a unique order-preserving bijection f between pairs of ordinals and ordinals by Proposition A.3.2.

Claim. If κ is an infinite cardinal, then f maps $\kappa \times \kappa$ to κ, and hence $\kappa \cdot \kappa = \kappa$. The proof of the claim is by induction on κ. For $\alpha, \beta \in \kappa$ let $P_{\alpha,\beta}$ be the set of predecessors of (α, β). Note that:

- $P_{\alpha,\beta}$ is contained in $\delta \times \delta$ with $\delta = \max(\alpha, \beta) + 1$.
- Since κ is infinite and $\alpha, \beta < \kappa$, the cardinality of δ is smaller than κ.
- By inductive assumption $|P_{\alpha,\beta}| \leq |\delta \times \delta| = |\delta| \cdot |\delta| \overset{\text{IA}}{=} |\delta| < \kappa$.

Hence, $f(\alpha, \beta) < \kappa$ since f is an order isomorphism and thus $f(\alpha, \beta) \in \kappa$.

Now (2) and (3) are simple consequences. Let $\mu := \max(\kappa, \lambda)$.

$$\mu \leq \kappa + \lambda \leq \mu + \mu \leq 2 \cdot \mu \leq \mu \cdot \mu = \mu$$
$$2^\kappa \leq \kappa^\kappa \leq (2^\kappa)^\kappa = 2^{\kappa \cdot \kappa} = 2^\kappa$$

□

COROLLARY A.4.3. *For a non-empty set A, the set $A^{<\omega} := \bigcup_{n\in\omega} A^n$ has cardinality* $\max(|A|, \aleph_0)$.

PROOF. Clearly, $|A| \leq |A^{<\omega}|$ and $\aleph_0 \leq |A^{<\omega}|$. On the other hand,

$$|A^{<\omega}| = \sum_{n\in\omega} |A|^n \leq (\sup_{n\in\omega} |A|^n) \cdot \aleph_0 = \max(|A|, \aleph_0).$$

The final equality holds because clearly $\sup_{n\in\omega} |A|^n = 1$ if $|A| = 1$, $\sup_{n\in\omega} |A|^n = \aleph_0$ if $2 \leq |A| \leq \aleph_0$, and $|A|$ if $|A| \geq \aleph_0$ by Theorem A.4.2 (1). □

The *Continuum Hypothesis (CH)* states that $\aleph_1 = 2^{\aleph_0}$, that is: there is no cardinal lying strictly between ω and the cardinality $|\mathbb{R}|$ of the continuum. The *Generalised Continuum Hypothesis (GCH)* states that $\kappa^+ = 2^\kappa$ for all infinite cardinals κ. As with CH, the GCH is known to be independent of ZFC, that is, there are models of ZFC where GCH is true , and models of ZFC where GCH is false (assuming that ZFC is consistent; see **[20]**).

Let A be a set that is linearly ordered by $<$. A subset $B \subseteq A$ is called *cofinal* if for every $a \in A$ there is some $b \in B$ with $a \leq b$. Any linear order contains a well-ordered cofinal subset.

Definition A.4.4. The *cofinality* $\mathrm{cf}(A)$ of A is the smallest ordinal that is order-isomorphic to a well-ordered cofinal subset of A.

Examples:
- If A has a greatest element, then the cofinality is one since the set consisting only of the greatest element is cofinal and must be contained in any other cofinal subset of A.
- A subset S of $\mathbb{N}$ is cofinal if and only if S is infinite, and thus $\mathrm{cf}(\omega) = \omega$.
- $\mathrm{cf}(2^\omega) > \omega$; see Exercise 94.

Lemma A.4.5. *For linearly ordered A we have* $\mathrm{cf}(\mathrm{cf}(A)) = \mathrm{cf}(A)$.

Proof. If $\{a_\alpha \mid \alpha < \beta\}$ is cofinal in A and $\{\alpha(\nu) \mid \nu < \mu\}$ is cofinal in β, then $\{a_{\alpha(\nu)} \mid \nu < \mu\}$ is cofinal in A. $\square$

Lemma A.4.6. *For linearly ordered A we have that $\mathrm{cf}(A)$ is a cardinal.*

Proof. Suppose for contradiction that $\mathrm{cf}(A)$ is not a cardinal. Choose a surjective map f from $|\mathrm{cf}(A)|$ to $\mathrm{cf}(A)$. This maps provides a cofinal sequence in $\mathrm{cf}(A)$ of length $|\mathrm{cf}(A)|$, and therefore $\mathrm{cf}(\mathrm{cf}(A)) \leq |\mathrm{cf}(A)| < \mathrm{cf}(A)$. This is in contradiction to $\mathrm{cf}(\mathrm{cf}(X)) = \mathrm{cf}(X)$ from Lemma A.4.5. $\square$

A cardinal κ is *regular* if $\mathrm{cf}(\kappa) = \kappa$, and *singular* otherwise. As we have seen above, $\aleph_0$ is an example of a regular cardinal. Lemma A.4.6 and Lemma A.4.5 show that $\mathrm{cf}(A)$ is a regular cardinal. Assuming the axiom of choice, we also have the following.

Proposition A.4.7. *Successor cardinals $\aleph_{\alpha+1}$ are regular.*

Proof. Suppose for contradiction that $\mathrm{cf}(\aleph_{\alpha+1}) \leq \aleph_\alpha$. Then $\aleph_{\alpha+1}$ would be the union of at most $\aleph_\alpha$ sets of cardinality at most $\aleph_\alpha$, contradicting item (1) in Theorem A.4.2. $\square$

Exercises.

(93) Show that $\aleph_\omega$ is singular.

(94) Show that $\mathrm{cf}(2^\omega) > \omega$.
Hint: write 2^ω as $\sum_{\nu<\mu} \kappa_\nu$ for $\mu := \mathrm{cf}(2^\omega)$, and apply König's theorem (Theorem A.4.1) with $(\kappa_\nu)_{\nu<\mu}$ and $(\lambda_\nu)_{\nu<\mu}$ where $\lambda_\nu := 2^\omega$ for all $\nu < \mu$.

APPENDIX B

Rings and Fields

In this section we collect some standard material from classical algebra that is typically taught in the second year; we freely follow Lang [**21**] and Dummit and Foote [**10**]. My impression is that Dummit and Foote is the more elementary and detailed presentation; on the other hand, the proof of Steinitz' theorem (Corollary B.4.12) is omitted there, but can be found in Lang's book, which is also the reference of Tent and Ziegler.

A commutative ring is a structure $\underline{R}$ over the signature $\tau_{\text{Ring}} = \{+, -, 0, 1, \cdot\}$ such that the reduct $(R; +, -, 0)$ of $\underline{R}$ is an Abelian group and such that the following axioms hold:

$$\begin{aligned}
&\forall x, y, z.\ (xy)z = x(yz) \\
&\forall x.\ 1x = x \\
&\forall x, y.\ xy = yx \\
&\forall x, y, z.\ x(y + z) = xy + xz
\end{aligned}$$

Let $\underline{R}$ be a ring. A nonzero element a of $\underline{R}$ is called a *zero divisor* if there is an $b \in R \setminus \{0\}$ such that $ab = 0$. An element $u \in R$ is called a *unit* if there is some $v \in R$ such that $uv = vu = 1$. In this terminology, a field is a commutative ring such that $0 \neq 1$ in which every nonzero element is a unit. Observe that a zero divisor cannot be a unit.

Lemma B.0.1. *Let $\underline{R}$ be a ring and $a, b, c \in R$ be such that a is not a zero divisor. If $ab = ac$ then $a = 0$ or $b = c$.*

Proof. If $ab = ac$ then $a(b - c) = 0$ so $a = 0$ or $b - c = 0$. □

B.1. Integral Domains

A commutative ring $\underline{R}$ with more than one element and without nonzero zero divisors is called an *integral domain*. An equivalent formulation is that $\underline{R}$ is a ring for which the set of nonzero elements is a commutative monoid with respect to multiplication. Examples of integral domains are fields, the ring of all integers, the p-adic integers, and rings of polynomials over integral domains.

A fundamental property of integral domains is that every subring of a field is an integral domain, and that, conversely, given any integral domain, one may construct a field that contains it as a subring, e.g., the *field of fractions*, introduced below. Hence, a ring is an integral domain if and only if it is isomorphic to a subring of a field.

Let $\underline{R}$ be an integral domain. Let $\sim$ be the equivalence relation defined on $A \times S$ by setting $(a, s) \sim (a', s')$ if and only if $s'a = sa'$. We write $\frac{a}{s}$ for the equivalence class containing (a, s).

Definition B.1.1. The *field of fractions* of an integral domain $\underline{R}$ is the field $\underline{Q}$ with the domain $Q := (R \times R \setminus \{0\})/_\sim$ where multiplication is defined by the rule

$$\frac{a}{s}\frac{a'}{s'} := \frac{aa'}{ss'}.$$

It is easy to verify that this is well defined, and that multiplication has a unit, namely $\frac{1}{1}$, and is associative. Addition is defined by the rule

$$\frac{a}{s} + \frac{a'}{s'} := \frac{s'a + sa'}{ss'}.$$

It is easy to verify that this is well defined, with neutral element $\frac{0}{1}$, and that addition and multiplication yield a ring structure.

Sometimes, the term *quotient field* is used, but we try to avoid this to avoid confusion with the quotients that will be introduced in the next section.

Note that the map $f_Q \colon R \to Q$ given by $f(a) := \frac{a}{1}$ is an (injective) ring homomorphism from $\underline{R}$ to $\underline{Q}$. Every element of $f_Q(R \setminus \{0\})$ is invertible in $\underline{Q}$: the inverse of $\frac{s}{1}$ is $\frac{1}{s}$. We often identify the elements of $\underline{R}$ with elements of $\underline{Q}$ along f_Q, and hence view $\underline{R}$ as a subring of $\underline{Q}$.

PROPOSITION B.1.2. *Let $\underline{R}$ and $\underline{B}$ be rings and $h \colon \underline{R} \to \underline{B}$ be a homomorphism such that all elements of $h(R)$ are invertible in $\underline{B}$. Let $\underline{Q}$ be the field of fractions of $\underline{R}$. Then there exists a unique homomorphism u from $\underline{Q}$ to $\underline{B}$ such that $u \circ f_Q = h$.*

EXAMPLE 81. If $\underline{R}$ is an integral domain, then the ring $\underline{R}[x_1, \dots, x_n]$ of polynomials over $\underline{R}$ is also an integral domain: this comes from the fact that if n and m are the degrees of the polynomials $f \in \underline{R}[x]$ and $g \in \underline{R}[x]$, then the degree of $fg \in \underline{R}[x]$ is $n + m$, so there are no nonzero zero divisors. If $\underline{K}$ is the field of fractions of $\underline{R}$, then the field of fractions of $\underline{R}[x_1, \dots, x_n]$ is denoted by $\underline{K}(x_1, \dots, x_n)$; its elements are called *rational functions*. A rational function can be written as $\frac{f(\bar{x})}{g(\bar{x})}$ where f and g are polynomials. △

B.2. Ideals and Quotient Rings

Let $\underline{R}$ be a commutative ring. A subset I of R is an *ideal* of $\underline{R}$ if

- $rI = Ir = I$ for all $r \in R$, and
- $(I; +)$ is a subgroup of $(R; +)$.

(i.e., the ideals of $\underline{R}$ are the domains of submodules of $\underline{R}$). For every ring homomorphism $h \colon \underline{R} \to \underline{S}$, the set $\{r \in R \mid h(r) = 0\}$, also called the *kernel* of the ring homomorphism h, is an ideal. Conversely, any ideal I of $\underline{R}$ gives rise to a congruence of $\underline{R}$, namely $C = \{(r_1, r_2) \mid r_1 - r_2 \in I\}$. We also write $\underline{R}/I$ for the structure $\underline{R}/C$, which will again be a ring since the ring axioms are universal. Note that each congruence class A can be written as $A = u + I$ for some $u \in A$. The homomorphism $h \colon \underline{R} \to \underline{R}/I$ that sends r to $r + I$ is called the *canonical homomorphism* from $\underline{R}$ to $\underline{R}/I$.

DEFINITION B.2.1 (principal ideal). For any $a \in R$, the set $Ra = aR$ is an ideal of $\underline{R}$, denoted by (a), the *principal ideal of $\underline{R}$ generated by a.*

Examples of principal ideals are $(0) = \{0\}$ (called the *trivial* or *zero ideal*) and $(1) = R$. Ideals different from R are called *proper ideals.*

EXAMPLE 82. Let $\underline{K}$ be a field and $\underline{K}[x]$ be the polynomial ring. Then the ideal (x) is the set of all $f \in \underline{K}[x]$ with zero constant coefficient. Let h be the canonical homomorphism from $\underline{K}[x]$ to $\underline{K}[x]/(x)$ given by $h(f(x)) := f(x) + (x)$. Then the restriction of h to $\underline{K}$ is an isomorphism between $\underline{K}$ and $\underline{K}[x]/(x)$: every element of $\underline{K}[x]/(x)$ can be written as $u + (x)$ where $u \in K$, and $h(u) = u + (x)$. △

PROPOSITION B.2.2. *A commutative ring $\underline{R}$ is a field if and only if its only ideals are $\{0\}$ and $\underline{R}$.*

Proof. $\underline{R}$ is a field if and only if every nonzero element is a unit. If $\underline{R}$ is a field then every nonzero ideal I contains a unit u with inverse v. Then for every $r \in R$

$$r = r(vu) = (rv)u \in I$$

hence $R = I$. Conversely, if $\{0\}$ and R are the only ideals of $\underline{R}$, then let $u \in R \setminus \{0\}$. By assumption $(u) = R$ and so $1 \in (u)$. Thus, there is some $v \in R$ such that $1 = uv$, so u is a unit. □

Note that for any finite family $\{I_j\}_{j \in A}$ of ideals of $\underline{R}$ the sum $\sum_{j \in A} I_j$ is an ideal of $\underline{R}$. If $A \subseteq R$, then $\sum_{a \in A}(a)$ is the smallest ideal of $\underline{R}$ that contains A. It is called the *ideal generated by* A. If $A = \{a_1, \dots, a_n\}$ then we write $(a_1, \dots, a_n)$ for the ideal generated by A.

Definition B.2.3. An integral domain $\underline{R}$ is called a *principal ideal domain (PID)* if every ideal of $\underline{R}$ is principal.

See Figure B.1 for an overview of the relationship with other fundamental concepts for rings and fields, along with examples that show that there is a chain of strict inclusions of the shown concepts.

Example 83. $\mathbb{Z}$ is a principal ideal domain: for if $I \subseteq \mathbb{Z}$ is a non-trivial ideal, then $I = n\mathbb{Z} = (n)$ for some $n \in \mathbb{N}$. △

Example 84. When $\underline{K}$ is a field, then the ring of polynomials $\underline{K}[x]$ is a principal ideal domain. △

Example 85. $\mathbb{Z}[x]$ is not a principal ideal domain: e.g. the ideal $(2, x)$ is not principal. Assume for contradiction that there exists $a(x) \in \mathbb{Z}[x]$ such that $(2, x) = (a(x))$. Since $2 \in (a(x))$ there must be some $p(x)$ such that $2 = p(x)a(x)$. The degree of $p(x)a(x)$ equals the degree of $p(x)$ plus the degree of $a(x)$, hence both $p(x)$ and $a(x)$ must be constant polynomials, i.e., integers. Since 2 is prime, $a(x), p(x) \in \{-2, -1, 1, 2\}$. If $a(x) \in \{-1, 1\}$ then every polynomial would be a multiple of $a(x)$ contrary to $(a(x)) \neq R$. If $a(x) \in \{-2, 2\}$ then $x \in (a(x)) = (2) = (-2)$ and so $x = 2q(x)$ for some $q \in \mathbb{Z}[x]$, clearly impossible. △

Example 86. $\mathbb{Q}[x, y]$ is not a principal ideal domain. This follows from Corollary B.2.10 below, because if it was, then $\mathbb{Q}[x]$ would have to be a field, which it is not. △

B.2.1. Prime Ideals. An ideal I of $\underline{R}$ is called *prime* if $I \neq R$ such that whenever $xy \in I$, then $x \in I$ or $y \in I$.

Proposition B.2.4. *An ideal I of $\underline{R}$ is prime if and only if $\underline{R}/I$ is an integral domain.*

Proof. An ideal I of $\underline{R}$ is prime if and only if $I \neq R$ and if $ab \in I$ then $a \in I$ or $b \in I$. This is the case if and only if $\underline{R}/I \neq \{I\}$ and if $abI = I$ then $aI = I$ or $bI = I$, i.e., if and only if $\underline{R}/I$ is an integral domain. □

Let $\underline{R}$ be a ring with a 1. Then the map $h \colon \mathbb{Z} \to R$ given by

$$h(n) := \underbrace{1 + \dots + 1}_{n \text{ times}}$$

is a ring homomorphism, and its kernel is an ideal $(n) = n\mathbb{Z}$, generated by $n \in \mathbb{N}$. If (n) is a prime ideal then $n = 0$ or $n = p$ for some prime number p. In the first case, $\underline{R}$ has a subring which is isomorphic to $\mathbb{Z}$; in this case, we say that $\underline{R}$ has *characteristic* 0. If on the other hand $n = p$ then we say that $\underline{R}$ has *characteristic* p, and $\underline{R}$ has a

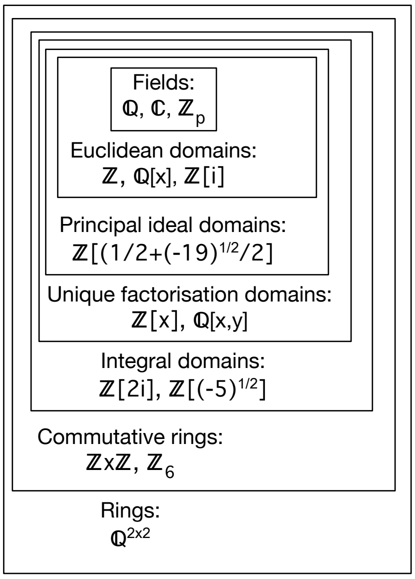

FIGURE B.1. From rings to fields: algebra reminder.

subring isomorphic to $\mathbb{Z}/n\mathbb{Z}$. If $\underline{K}$ is a field, then $\underline{K}$ has characteristic 0 or $p > 0$. In the first case, it contains an isomorphic copy of $\mathbb{Q}$, and in the second case, it contains an isomorphic copy of $\mathbb{F}_q$. In either case, this subfield will be called the *prime field* (contained in $\underline{K}$). The prime field is the smallest subfield of $\underline{K}$ containing 1, and it has no non-trivial automorphisms.

DEFINITION B.2.5. Let r be a nonzero non-unit element of an integral domain $\underline{R}$. Then r is called

- *irreducible* in $\underline{R}$ if whenever $r = ab$ for $a, b \in R$, then at least one of a or b must be a unit in $\underline{R}$;
- *prime* in $\underline{R}$ if the ideal (r) is prime.

PROPOSITION B.2.6. *Every prime element in an integral domain is irreducible.*

PROOF. Suppose that (p) is a nonzero prime ideal and $p = ab$. Then $ab = p \in (p)$, so by definition of prime ideals one of a or b, say a, is in (p). Thus $a = pr$ for some r. This implies that $p = ab = prb$, so $rb = 1$ (Lemma B.0.1) and b is a unit. This shows that p is irreducible. □

In Proposition B.2.11 we will see a converse of Proposition B.2.6 in principal ideal domains.

B.2.2. Maximal Ideals. An ideal I of $\underline{R}$ is called *maximal* if $I \neq R$ and if there is no ideal $J \neq R$ strictly containing I.

PROPOSITION B.2.7. *An ideal I of a commutative ring $\underline{R}$ is maximal if and only if $\underline{R}/I$ is a field.*

PROOF. By Proposition B.2.2, $\underline{R}/I$ is a field if and only if $\{0\}$ is the only proper ideal of $\underline{R}/I$. Note that J is an ideal of $\underline{R}/I$ if and only if $\bigcup J$ is an ideal of $\underline{R}$ that contains I; so I is maximal if and only if $\{0\}$ is the only proper ideal of $\underline{R}/I$. □

COROLLARY B.2.8. *Every maximal ideal of a commutative ring $\underline{R}$ is prime.*

Proof. If I is a maximal ideal, then $\underline{R}/I$ is a field by Proposition B.2.7, which is in particular an integral domain. Proposition B.2.4 the implies that I is prime. $\square$

In principal ideal domains, we have a converse for Corollary B.2.8.

Proposition B.2.9. *Every nonzero prime ideal in a principal ideal domain $\underline{R}$ is maximal.*

Proof. Let (p) be a nonzero prime ideal in $\underline{R}$ and let $I = (m)$ be any ideal containing (p). We must show that $I = (p)$ or $I = R$. Since $p \in I$ there is an $r \in R$ such that $p = rm$. Since (p) is prime and $p = rm \in (p)$ either $r \in (p)$ or $m \in (p)$. If $m \in (p)$ then $(p) = (m) = I$. If $r \in (p)$ write $r = sp$. In this case $p = rm = psm$. Since $\underline{R}$ is in particular an integral domain, $p \neq 0$ is not a zero divisor, and hence $sm = 1$ (Lemma B.0.1). Hence m is a unit, so $I = R$. $\square$

Corollary B.2.10. *Let $\underline{R}$ be a commutative ring such that $\underline{R}[x]$ is a principal ideal domain. Then $\underline{R}$ must be a field.*

Proof. Since $\underline{R}[x]$ is by assumption in particular an integral domain, so is its subring $\underline{R}$. By Proposition B.2.4, the non-zero ideal (x) in $\underline{R}[x]$ is prime because $\underline{R}[x]/(x)$ is isomorphic to the integral domain $\underline{R}$ (see Example 82). By Proposition B.2.9, (x) is maximal, hence $\underline{R}[x]/(x)$ is a field by Proposition B.2.7. $\square$

Note that this shows in particular that $\mathbb{Q}[x, y]$ is not a principal ideal domain: we have $\mathbb{Q}[x, y] = \underline{R}[x]$ for $\underline{R} := \mathbb{Q}[y]$, but $\underline{R}$ is not a field. Now comes the promised converse of Proposition B.2.6 for principal ideal domains.

Proposition B.2.11. *Let $\underline{R}$ be a principal ideal domain and let $I = (p)$ be a nonzero ideal. Then the following are equivalent.*

(1) p is prime;
(2) p is irreducible;
(3) I is maximal;
(4) I is prime.

Proof. (1) $\Rightarrow$ (2) follows from Proposition B.2.6.

For (2) $\Rightarrow$ (3), let p be irreducible. If I is any ideal containing (p) then by assumption I is principal, i.e., $I = (m)$ for some $m \in R$. Since $p \in (m)$ there exists $r \in R$ such that $p = rm$. But p is irreducible, so by definition either r or m is a unit. If r is a unit then $(p) = (m)$, and if m is a unit then $(m) = R$. It follows that I is maximal.

(3) $\Rightarrow$ (4) follows from Corollary B.2.8.

(4) $\Rightarrow$ (1) holds by definition. $\square$

Exercises.

(95) A *unique factorisation domain* is an integral domain in which every nonzero non-unit element can be written as a product of irreducible elements, uniquely up to order and units. Show that every principal ideal domain is a unique factorisation domain. See Figure B.1.

(96) Show that every principal ideal domain is a *Dedekind domain*, i.e., an integral domain in which every nonzero proper ideal factors into a product of prime ideals.

(97) Show that $\mathbb{Q}[x, y]$ is an example of a unique factorisation domain which is not a Dedekind domain.

(98) Show that a Dedekind domain is a unique factorisation domain if and only if it is a principal ideal domain.

B.3. Polynomial Rings

Let $\underline{R}$ be a commutative ring. We write $\underline{R}[x]$ for the ring of polynomials over $\underline{R}$. Let $\underline{S}$ be a subring of $\underline{R}$. If $p(x) \in \underline{S}[x]$, we write $p^{\underline{R}}\colon R \to R$ for the corresponding polynomial function. Note that $p \mapsto p^{\underline{R}}(0)$ is a ring homomorphism from $\underline{S}[x]$ to $\underline{R}$.

Let $b \in R$. We write $\underline{S}[b]$ for the subring of $\underline{R}$ generated by $S \cup \{b\}$ (i.e., the smallest substructure of $\underline{R}$ that contains $S \cup \{b\}$); this is also called a *ring adjunction*. Note that the elements of $\underline{S}[b]$ are precisely

$$\{p^{\underline{R}}(b) \mid p \in \underline{S}[x]\}.$$

If the map $p \mapsto p^{\underline{R}}(b)$ is an isomorphism between $\underline{S}[x]$ and $\underline{S}[b]$ then we say that x is *transcendental* over $\underline{S}$.

We have already mentioned that for any field $\underline{K}$, the ring of polynomials $\underline{K}[x]$ is a principal ideal domain. A polynomial is called *monic* if the leading coefficient is 1.

PROPOSITION B.3.1. *Let $\underline{K}$ be a field and $p \in \underline{K}[x]$ of degree $n \geq 0$. Then f has at most n roots in $\underline{K}$, and if a is a root of f in $\underline{K}$, then $x - a$ divides $f(x)$.*

PROOF. By polynomial division and induction. □

PROPOSITION B.3.2. *Let $p \in \underline{K}[x]$ be irreducible. Then the quotient $\underline{K}[x]/(p(x))$ is a field.*

PROOF. Since $p(x)$ is irreducible, the ideal $(p(x))$ is maximal by Proposition B.2.11, so $\underline{K}[x]/(p(x))$ is a field by Proposition B.2.7. □

B.4. Field Extensions

A *field extension* $\underline{E}$ of a field $\underline{F}$ is a field that extends $\underline{F}$ in the sense of Section 1.1.3. We may view $\underline{F}$ as a vector space over $\underline{E}$, and write $[\underline{E} : \underline{F}]$ for the dimension of the vector space, called the *degree* of the extension. We say that $\underline{E}$ is a *finite extension* or an *infinite extension* of $\underline{F}$ according to whether the degree is finite or not. If $S \subseteq E$ then there exists a smallest subfield of $\underline{E}$ that contains F and S, denoted by $\underline{F}(S)$, and called the field *generated* by S over $\underline{F}$ (this is not the same as the notation $\underline{F}[S]$ from Section B.3, since $\underline{F}(S)$ must also contain the multiplicative inverses for non-zero elements). When $S = \{a_1, \dots, a_n\}$ is finite one writes $\underline{F}(a_1, \dots, a_n)$ instead of $\underline{F}(\{a_1, \dots, a_n\})$. If S consists of a single element a, then the extension $\underline{F}(a)$ of $\underline{F}$ is called *simple* and a is called a *primitive element* of the extension. Note that the elements of $\underline{F}(a)$ are precisely the set of elements that can be written as $p(a)/q(a)$ where $p, q \in \underline{F}[x]$ (for $q(a) \neq 0$).

B.4.1. Algebraic Field Extensions. An element a of a field extension $\underline{E}$ of $\underline{F}$ is called *algebraic over* $\underline{F}$ if a is the root of some non-zero polynomial $f(x)$ with coefficients from $\underline{F}$.

LEMMA B.4.1. *Let $\underline{E}$ be a field extension of $\underline{F}$. Then $a \in E$ is algebraic over $\underline{F}$ if and only if a is not transcendental over $\underline{F}$ (see Section B.3).*

PROOF. If a is not transcendental, then $p \mapsto p^{\underline{E}}(a)$ is not injective, i.e., there are distinct $p_1, p_2 \in \underline{F}[x]$ such that $p_1^{\underline{E}}(a) = p_2^{\underline{E}}(a)$. Hence, a is a zero of the non-zero polynomial $p_1 - p_2$, and algebraic over $\underline{F}$.

Conversely, if a is a zero of the non-zero polynomial $p \in \underline{F}[x]$, then $p^{\underline{E}}(a) = 0$, but the map $p \mapsto p^{\underline{E}}(a)$ sends the 0 polynomial in $\underline{F}[x]$ to 0, too, so a is not transzendental. □

LEMMA B.4.2. *If a is transcendental over $\underline{F}$, then the isomorphism $p \mapsto p^{\underline{E}}(a)$ between $\underline{F}[x]$ and $\underline{F}[a]$ can be extended to an isomorphism between the field of rational functions $\underline{F}(x)$ (Example 81) and $\underline{F}(a)$.*

PROOF. If $f, g \in \underline{F}[x]$, and g is non-zero, then $g(a) \neq 0$, and we map f/g to $f^{\underline{E}}(a)/g^{\underline{E}}(a) \in \underline{F}(a)$. This map is clearly a homomorphism, and it is injective: suppose that $f^{\underline{E}}(a)/g^{\underline{E}}(a) = (f')^{\underline{E}}(a)/(g')^{\underline{E}}(a)$ for $f', g' \in \underline{F}[x]$, g' non-zero, then $f(a)g'(a) - f'(a)g(a) = 0$. By Lemma B.4.1, a is not algebraic, so $fg' - f'g$ must be the zero-polynomial, which in turn means that $f/g = f'/g'$. $\square$

LEMMA B.4.3. *If $\underline{E}$ is a field extension of $\underline{F}$ and $a \in E$ is algebraic over $\underline{F}$, then there exists an irreducible $p \in \underline{F}[x]$ such that $p(a) = 0$. Furthermore, if $f \in \underline{F}[x]$ and $f(a) = 0$, then p divides f.*

PROOF. There exists a polynomial $p \in \underline{F}[x]$ such that $p(a) = 0$; choose p of minimal degree. We claim that p is irreducible in $\underline{F}[x]$. For if $p = tq$ for $t, q \in \underline{F}[x]$ then $0 = p(a) = t(a)q(a)$, and so a would be a root of at least one of t and q, say of t. Hence, $\deg(t) = \deg(p)$, since p is of minimal degree among those polynomials with root a. Hence, q must be a constant polynomial, as required. To show that p divides every $f \in \underline{F}[x]$ with root a, we write $f = pq + r$ where $q, r \in \underline{F}[x]$ and $\deg(r) < \deg(p)$. Then $0 = f(a) = p(a)q(a) + r(a) = r(a)$, so a is a root of r. Hence $r = 0$, and p divides f. $\square$

The proof of Lemma B.4.3 shows that there exists a unique monic polynomial in $\underline{F}[x]$ that is irreducible and has a as a root, and which is called the *minimal polynomial of a over $\underline{F}$*.

PROPOSITION B.4.4. *Let $\underline{E}$ be a field extension of $\underline{F}$. Then $a \in E$ is algebraic over $\underline{F}$ if and only if $\underline{F}(a)$ is a finite extension of $\underline{F}$.*

PROOF. If a is algebraic then the degree of $\underline{F}(a)$ over $\underline{F}$ is the degree of the minimal polynomial for a over $\underline{F}$. Conversely, suppose that $\underline{F}(a)$ has degree n over $\underline{F}$. Then $1, a, a^2, \ldots, a^n$ cannot be linearly independent over $\underline{F}$. So there are $b_0, b_1, \ldots, b_n \in F$, not all zero, such that $b_0 + b_1 a + \cdots + b_n a^n = 0$. Hence, a is the root of the non-zero polynomial $b_0 + b_1 x + \cdots + b_n x^n$. $\square$

A field extension $\underline{E}$ of $\underline{F}$ is *algebraic* if every element of $\underline{E}$ is algebraic over $\underline{F}$. Field extensions that are not algebraic are called *transcendental*.

COROLLARY B.4.5. *Every finite field extension $\underline{E}$ of $\underline{F}$ is algebraic over $\underline{F}$.*

PROOF. If $a \in E$, then $\underline{F}(a)$ is a subspace of the vector space $\underline{E}$ over $\underline{F}$. Hence $[\underline{F}(a) : \underline{F}] \leq [\underline{E} : \underline{F}]$ and so a is algebraic over $\underline{F}$ by Proposition B.4.4. $\square$

A literal converse of Lemma B.4.5 is false: for example the subfield of $\mathbb{C}$ consisting of all algebraic numbers is an infinite algebraic extension of $\mathbb{Q}$. The following lemma prepares an exact characterisation of finite field extensions (Proposition B.4.7 below).

LEMMA B.4.6. *Let $\underline{F}$ be a subfield of $\underline{K}$, and $\underline{K}$ a subfield of $\underline{L}$. Then $[\underline{L} : \underline{F}] = [\underline{L} : \underline{K}][\underline{K} : \underline{F}]$.*

PROOF. Suppose first that $[\underline{L} : \underline{K}] = m$ and $[\underline{K} : \underline{F}] = n$ are finite. Let $a_1, \ldots, a_m$ be a basis for $\underline{L}$ over $\underline{K}$ and let $b_1, \ldots, b_n$ be a basis for $\underline{K}$ over $\underline{F}$. Then every $c \in \underline{L}$ can be written as a linear combination

$$u_1 a_1 + \cdots + u_m a_m$$

where $u_1, \dots, u_m \in K$. Each u_i, for $i \in \{1, \dots, m\}$ can be written as a linear combination

$$u_i = v_{i,1} b_1 + \cdots + v_{i,n} b_n \tag{3}$$

where the $v_{i,j} \in F$. Substituting these expressions we obtain the linear combination

$$c = \sum_{i \in \{1,\dots,m\}} \sum_{j \in \{1,\dots,n\}} v_{i,j} a_i b_j$$

so the elements $a_i b_j$ span $\underline{L}$ as a vector space over $\underline{F}$.

Suppose now that

$$\sum_{i \in \{1,\dots,m\}} \sum_{j \in \{1,\dots,n\}} v_{i,j} a_i b_j = 0$$

for some coefficients $v_{ij} \in F$. Then defining the elements $u_i \in K$ by equation 3 above this equation could be rewritten to $u_1 a_1 + \cdots + u_m a_m = 0$, and since $a_1, \dots, a_m$ form a basis for $\underline{L}$ over $\underline{K}$ this implies that $u_1 = \cdots = u_m = 0$. So we have $v_{i,1} b_1 + \cdots + v_{i,n} b_n = 0$, and since $b_1, \dots, b_n$ for a basis for $\underline{K}$ over $\underline{F}$ this implies that $v_{i,j} = 0$ for all $i \le m$ and $j \le n$. Hence the $a_i b_j$ are linearly independent over $\underline{F}$, so they form a basis for $\underline{L}$ over $\underline{F}$ and $[\underline{L} : \underline{F}] = mn$ as claimed.

If one of $[\underline{L} : \underline{K}]$ or $[\underline{K} : \underline{F}]$ is infinite, then $[\underline{L} : \underline{F}]$ is infinite, too. □

PROPOSITION B.4.7. *$\underline{E}$ is a finite extension of $\underline{F}$ if and only if $\underline{E}$ is generated by a finite number of elements of E that are algebraic over $\underline{F}$.*

PROOF. If $[\underline{E} : \underline{F}] = n \in \mathbb{N}$, let $a_1, \dots, a_n$ be a basis for $\underline{E}$ as a vector space over $\underline{F}$. By Lemma B.4.6, $[\underline{F}(a_i) : \underline{F}]$ divides $[\underline{E} : \underline{F}] = n$ for $i \in \{1, \dots, n\}$, so Proposition B.4.4 implies that each a_i is algebraic over $\underline{F}$.

For the converse, suppose that $\underline{E} = \underline{F}(a_1, \dots, a_n)$ where each of $a_1, \dots, a_n$ is algebraic over $\underline{F}$. We have $\underline{F}(a_1, \dots, a_n) = \underline{F}(a_1) \dots (a_n)$, and by Proposition B.4.4, and $[\underline{F}(a_1) : \underline{F}]$ is finite by Proposition B.4.4. note that $[\underline{F}(a_1)(a_2) : \underline{F}(a_1)]$ is equal to the degree of the minimal polynomial of a_2 over $\underline{F}(a_1)$, which is at most the degree of the minimal polynomial p of a_2 over $\underline{F}$ (the degrees are equal if p remains irreducible over $\underline{F}(a_1)$). By Lemma B.4.6 and induction it follows that $[\underline{E} : \underline{F}]$ is finite. □

LEMMA B.4.8. *Let $\underline{L}$ be an algebraic extension of $\underline{K}$, and $\underline{K}$ an algebraic extension of $\underline{F}$. Then $\underline{L}$ is an algebraic extension of $\underline{F}$.*

PROOF. Let $a \in L$. Then a satisfies an equation $b_n a^n + \cdots + b_0 = 0$ where $b_1, \dots, b_n \in K$, $b_0 \neq 0$. Consider the subfield $\underline{F}(a, b_0, b_1, \dots, b_n)$ of $\underline{L}$. Since $\underline{K}$ is algebraic over $\underline{F}$, the elements $b_0, \dots, b_n$ are algebraic over $\underline{F}$, so $[\underline{F}(b_0, \dots, b_n) : \underline{F}]$ is finite by Proposition B.4.7. Since a satisfies the equation above, we have that $[\underline{F}(a, b_0, \dots, b_n) : \underline{F}(b_0, \dots, b_n)] \le n$, and by Lemma B.4.6 it follows that $[\underline{F}(a, b_0, \dots, b_n) : \underline{F})]$ is finite, which proves that a is algebraic over $\underline{F}$. □

B.4.2. Algebraically Closed Field Extensions. A field $\underline{K}$ is called *algebraically closed* if it has no proper algebraic extension, i.e., it contains a root for every non-constant polynomial with coefficients in $\underline{K}$. The typical example of an algebraically closed field is the field of complex numbers. Also note that no finite field $\underline{K}$ is algebraically closed, because if $K = \{a_1, \dots, a_n\}$ then the polynomial $(x - a_0) \cdots (x - a_n) + 1$ has no root in $\underline{K}$.

PROPOSITION B.4.9. *Let $\underline{K}$ be a field and $p \in \underline{K}[x]$ of degree at least one. Then there exists an extension $\underline{E}$ of $\underline{K}$ in which p has a root.*

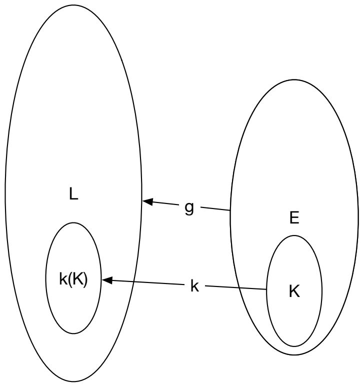

FIGURE B.2. Illustration for Theorem B.4.11.

PROOF. We may assume that p is irreducible. Then the quotient ring $\underline{E} := \underline{K}[x]/(p(x))$ is a field by Proposition B.3.2. Let $h\colon \underline{K}[x] \to \underline{K}[x]/(p(x))$ be the canonical homomorphism given by $f(x) \mapsto f(x) + (p(x))$, and let e be the restriction of h to $\underline{K}$. Then $e\colon \underline{K} \to \underline{E}$ is not constant 0 since it maps $1 \in K$ to $1 \in F$. Since $\{0\}$ is the only proper ideal of a field (Proposition B.2.2), e must be an embedding. We identify $\underline{K}$ with its image in $\underline{E}$ via e and view $\underline{K}$ as a subfield of $\underline{E}$. Then

$$\begin{aligned} p(e(x)) &= e(p(x)) && \text{(since } e \text{ is a homomorphism)} \\ &= p(x) + (p(x)) && \text{(an element of } \underline{E}\text{)} \\ &= 0^{\underline{E}} \end{aligned}$$

so $\underline{E}$ does indeed contain a root of the polynomial $p(x)$. Hence, $\underline{E}$ is a field extension of $\underline{K}$ with the desired property. □

Recall from Corollary 7.4.3 that every field has an algebraically closed field extension. Clearly, algebraically closed field extensions of a field are not unique: they might be of different cardinality. But we will prove that algebraically closed *algebraic* field extensions are unique up to isomorphism (Corollary B.4.12). First we prove existence of such field extensions.

COROLLARY B.4.10. *Every field $\underline{K}$ has an algebraically closed* algebraic *field extension.*

PROOF. By Corollary 7.4.3 there exists an algebraically closed field extension $\underline{E}$ of $\underline{K}$. Let $\underline{A}$ be the union of all subfields of $\underline{E}$ that are algebraic over $\underline{K}$; this is itself a subfield of $\underline{E}$, and certainly algebraic over $\underline{K}$. If $p \in \underline{A}[x]$ has degree at least one, then it has a root a in E, so a is algebraic over $\underline{A}$. If a is algebraic over $\underline{A}$, then a is already algebraic over $\underline{K}$ by Lemma B.4.8, and therefore $a \in A$. We have shown that $\underline{A}$ is algebraically closed. □

B.4.3. Algebraic Closure.

THEOREM B.4.11. *Let $\underline{K}$ be a field, $\underline{E}$ an algebraic extension of $\underline{K}$, and $k\colon \underline{K} \to \underline{L}$ an embedding of $\underline{K}$ into an algebraically closed field $\underline{L}$. Then there exists an extension of k to $g\colon \underline{E} \hookrightarrow \underline{L}$. If $\underline{E}$ is algebraically closed and $\underline{L}$ is algebraic over $k(K)$, then any such extension of k is an isomorphism between $\underline{E}$ and $\underline{L}$.*

PROOF. Let S be the set of all pairs $(\underline{F}, f)$ where $\underline{F}$ is a subfield of $\underline{E}$ containing $\underline{K}$ and f is an extension of k to an embedding of $\underline{F}$ into $\underline{L}$. We define the following partial order on S: if $(\underline{F}, f), (\underline{F}', f') \in S$ define $(\underline{F}, f) \leq (\underline{F}', f')$ if $F \subseteq F'$ and f is the restriction of f' to F. Note that $S \neq \emptyset$ because it contains $(\underline{K}, k)$. Also note that every chain $(\underline{F}_i, f_i)_{i \in I}$ of $(S, \leq)$ has an upper bound in S: define $F := \bigcup_{i \in I} F_i$ and define f on F to be equal to f_i on F_i for each $i \in I$. Using Zorn's Lemma (Theorem A.2.1), let $(\underline{G}, g)$ be a maximal element in S.

Then g is an extension of k, and we claim that $G = E$. Indeed, if $a \in E \setminus G$, then $\underline{G}(a)$ is algebraic over $\underline{G}$ because E is algebraic over $\underline{K}$. For any $b \in \underline{G}(a)$, let $p \in \underline{G}[x]$ be the minimal polynomial of b over $\underline{G}$. Note that $p(x) = t(x, c_1, \dots, c_n)$ where t is a term in the language of rings and $c_1, \dots, c_n$ are constant symbols for elements of G. Since $\underline{L}$ is algebraically closed, there exists a b' such that $\underline{L} \models (b', g(c_1), \dots, g(c_n))$, and we define $g(b) := b'$. Then g has an extension to $\underline{G}(a)$: if $b \in \underline{G}(a)$, we can write it in the form $s(a, d_1, \dots, d_m)$ for some term s in the language of rings and $d_1, \dots, d_m$ constant symbols for elements of G. Then the map $s(a, d_1, \dots, d_m) \mapsto s(b, g(d_1), \dots, g(d_m))$ is well-defined (independent from the choice of s) and gives a embedding of $\underline{G}(a)$ into $\underline{L}$ that extends g. This is contradicting the maximality of $(\underline{G}, g)$, and shows that g is an embedding $\underline{E} \hookrightarrow \underline{L}$.

If $\underline{E}$ is algebraically closed and $\underline{L}$ is algebraic over $g(K)$, then $g(E)$ is algebraically closed and $\underline{L}$ is algebraic over $g(E)$, hence $L = g(E)$. $\square$

Corollary B.4.12 (Steinitz). *Let $\underline{K}$ be a field and let $\underline{E}$ and $\underline{E}'$ be algebraic extensions of $\underline{K}$. Assume that $\underline{E}$ and $\underline{E}'$ are algebraically closed. Then there exists an isomorphism between $\underline{E}$ and $\underline{E}'$ that fixes the elements of $\underline{K}$.*

Proof. By Theorem B.4.11 the identity map on $\underline{K}$ can be extended to a map from $\underline{E}$ to $\underline{E}'$. Since $\underline{E}$ is algebraically closed and $\underline{E}$ is algebraic over $h(K)$, by the second part of Theorem B.4.11, the extension is an isomorphism between $\underline{E}$ and $\underline{E}'$. $\square$

Hence, algebraically closed algebraic extensions are uniquely determined up to isomorphism, and will be called the *algebraic closure* of $\underline{K}$.

Example 87. The subfield of $\mathbb{C}$ consisting of the complex numbers that are algebraic over $\mathbb{Q}$ is countable, and the algebraic closure of $\mathbb{Q}$. $\triangle$

B.4.4. Transcendence degree. The notion of *transcendence degree* is analogous to the notion of *dimension* in linear algebra.

Definition B.4.13. Let $\underline{E}$ be a field extension of a field $\underline{F}$.

- A subset $S \subseteq E$ is called *algebraically independent over* $\underline{F}$ if for every $p \in \underline{F}[x_1, \dots, x_n]$ and $a_1, \dots, a_n \in S$, if $p(a_1, \dots, a_n) = 0$ then p is the zero polynomial.
- A *transcendence base* (or *transcendence basis*) of $\underline{E}$ over $\underline{F}$ is an algebraically independent subset $S \subseteq E$ such that $\underline{E}$ is algebraic over $\underline{F}(S)$.
- The *transcendence degree* of $\underline{E}$ over $\underline{F}$ is the cardinality of a transcendence base.

Note that $a \in E$ is transcendent over $\underline{F}$ if and only if a is algebraically independent over $\underline{F}$ if and only if $\{a\}$ is a transcendence basis of $\underline{F}(a)$ over $\underline{F}$.

Example 88. $\{\sqrt{2}, e\}$ has transcendence degree one over $\mathbb{Q}$, since $\sqrt{2}$ is algebraic over $\mathbb{Q}$, but e is not. It is not known whether $\{\pi, e\}$ are algebraically independent over $\mathbb{Q}$. The field of complex numbers has transcendence degree 2^κ. $\triangle$

The following facts can be shown analogously to the corresponding statements about bases of vector spaces; these facts also follow from more general facts about pre-geometries that we present in Section 10.1.

Proposition B.4.14. *Let $\underline{E}$ be a field extension of $\underline{F}$. Then*

- $S \subseteq E$ *is a transcendence basis if and only if it is a* maximal *algebraically independent subset of* E*;*
- *Every field has a transcendence basis (this statement relies on Zorn's lemma)*
- *Any two transcendence bases for $\underline{E}$ over $\underline{F}$ have the same cardinality.*

APPENDIX C

Ramsey's theorem

We denote the set $\{0, \dots, n-1\}$ also by $[n]$. Subsets of a set of cardinality s will be called s-subsets in the following. Let $\binom{M}{s}$ denote the set of all s-subsets of M. We also refer to mappings $\chi \colon \binom{M}{s} \to [c]$ as a *coloring* of M (with the *colors* $[c]$). In Ramsey theory, one writes

$$L \to (m)^s_c$$

if for every $\chi \colon \binom{L}{s} \to [c]$ there exists an $M \subseteq L$ with $|M| = m$ such that χ is constant on $\binom{M}{s}$. In the following, ω denotes the cardinality of $\mathbb{N}$. Note the following.

- For all $n \in \mathbb{N}$ we have $[n+1] \to (2)^1_n$: this is the pigeon-hole principle.
- For all $c \in \mathbb{N}$ we have $\mathbb{N} \to (\omega)^1_c$: this is the infinite pigeon-hole principle.

We first state and prove a special case of Ramsey's theorem.

Theorem C.0.1. $\mathbb{N} \to (\omega)^2_2$.

This statement has the following interpretation in terms of undirected graphs: every countably infinite undirected graph either contains an infinite *clique* (a complete subgraph) or an infinite *independent set* (a subgraph without edges).

Proof. Let $\chi \colon \binom{\mathbb{N}}{2} \to [2]$ be a 2-colouring of the edges of $\binom{\mathbb{N}}{2}$. We define an infinite sequence $x_0, x_1, \dots$ of numbers from $\mathbb{N}$ and an infinite sequence $V_0 \supseteq V_1 \supseteq \cdots$ of infinite subsets of $\mathbb{N}$. Start with $V_0 := \mathbb{N}$ and $x_0 = 0$. By the infinite pigeon-hole principle, there is a $c_0 \in [2]$ such that $\{v \in V_0 \mid \chi(x_0, v) = c_0\} =: V_1$ is infinite. We now repeat this procedure with any $x_1 \in V_1$ and V_1 instead of V_0. Continuing like this, we obtain sequences $(c_i)_{i \in \mathbb{N}}$, $(x_i)_{i \in \mathbb{N}}$, $(V_i)_{i \in \mathbb{N}}$.

Again by the infinite pigeon-hole principle, there exists $c \in [2]$ such that $c_i = c$ for infinitely many $i \in \mathbb{N}$. Then $P := \{x_i \mid c_i = c\}$ has the desired property. To see this, let $i < j$ be such that $x_i, x_j \in P$. Then $\chi(\{x_i, x_j\}) = c_i = c$. $\square$

We now state Ramsey's theorem in it's full strength; the proof is similar to the proof of Theorem C.0.1 shown above.

Theorem C.0.2 (Ramsey's theorem). *Let $s, c \in \mathbb{N}$. Then $\mathbb{N} \to (\omega)^s_c$.*

A proof of Theorem C.0.2 can be found in **[19]** (Theorem 5.6.1); for a broader introduction to Ramsey theory see **[14]**. Compactness and Theorem C.0.2 imply the following finite version of Ramsey's theorem (exercise). In fact, full compactness is not needed, as we will see in our proof which is based on Kőnig's tree lemma.

A *walk* in a graph (V, E) (see Example 1) is a sequence $x_0, x_1, \dots, x_n \in V$ with the property that $(x_i, x_{i+1}) \in E$ for all $i \in \{1, \dots, n-1\}$. A walk is a *path* if all its vertices are distinct. A *cycle* is a walk of length at least three of the form $x_0, x_1, \dots, x_n = x_0$ such that $x_1, \dots, x_n$ are pairwise distinct. A *tree* is a connected graph (V, E) (see Section 1.1.6) without cycles. The *degree* of a vertex $u \in V$ is the number of vertices $v \in V$ such that $\{u, v\} \in E$.

LEMMA C.0.3 (Kőnig's Tree Lemma). *Let (V,E) be a tree such that every vertex in V has finite degree, and let $v_0 \in V$. If there are arbitrarily long paths that start in v_0, then there is an infinitely long path that starts in v_0.*

PROOF. Since the degree of v_0 is finite, there exists a neighbour v_1 of v_0 such that arbitrarily long paths start in v_0 and continue in v_1 (by the infinite pigeonhole principle). We now construct the infinitely long path by induction. Suppose we have already found a sequence $v_0, v_1, \dots, v_i$ that can be continued to arbitrarily long paths in (V,E). Since the degree of v_i is finite, v_i must have a neighbour v_{i+1} in $V \setminus \{v_0, v_1, \dots, v_i\}$ such that $v_0, v_1, \dots, v_{i+1}$ can be continued to arbitrarily long paths in (V,E). In this way, we define an infinitely long path $v_0, v_1, v_2, \dots$ in (V,E). $\square$

The degree assumption in Lemma C.0.3 is necessary (exercise).

THEOREM C.0.4 (Finite version of Ramsey's theorem). *For all $c, m, s \in \mathbb{N}$ there is an $l \in \mathbb{N}$ such that $[l] \to (m)^s_c$.*

PROOF. A proof by contradiction: suppose that there are positive integers c, m, s such that for all $l \in \mathbb{N}$ there is a $\chi \colon \binom{[l]}{s} \to [c]$ such that $(*)_{[l]}$ for all m-subsets M of $[l]$ the mapping χ is not constant on $\binom{M}{s}$. We construct a tree as follows. The vertices are the maps $\chi \colon \binom{[l]}{s} \to [c]$ that satisfy $(*)_{[l]}$. We make the vertex $\chi \colon \binom{[l]}{s} \to [c]$ adjacent to $\chi \colon \binom{[l+1]}{s} \to [c]$ if χ is a restriction of χ'. Clearly, every vertex in the tree has finite degree. By assumption, there are arbitrarily long paths that start in the vertex χ_0 where χ_0 is the map with the empty domain. By Kőnig's tree lemma the tree contains an infinite path $\chi_0, \chi_1, \dots$ We use this to define a map $\chi_{\mathbb{N}} \colon \binom{\mathbb{N}}{s} \to [c]$ as follows. For every $x \in \mathbb{N}$, there exists a $c_0 \in [c]$ and an $i_0 \in \mathbb{N}$ such that $\chi_i(x) = c_0$ for all $i \geq i_0$. Define $\chi_{\mathbb{N}}(x) := c_0$. Then $\chi_{\mathbb{N}}$ satisfies $(*)_{\mathbb{N}}$, a contradiction to Theorem C.0.2. $\square$

Exercises.

(99) Let $(X;<)$ be a partially ordered set on a countably infinite set X. Show that $(X;<)$ contains an infinite chain, or an infinite antichain.

(100) Show that an infinite sequence of elements of a totally ordered set contains one of the following:
- a constant subsequence;
- a strictly increasing subsequence;
- a strictly decreasing subsequence.

Derive the Bolzano-Weierstrass theorem (every bounded sequence in $\mathbb{R}^n$ has a convergent subsequence), using the completeness property of $\mathbb{R}^n$.

Ramsey's theorem can be used to prove the Erdös-Makkai theorem which we need in the proof of Theorem 8.1.5 about equivalent characterisations of stability.

THEOREM C.0.5 (Erdős-Makkai). *Let A be an infinite set and let $\mathcal{S}$ be a set of subsets of A with $|A| < |\mathcal{S}|$. Then there are a sequence $(a_i)_{i \in I}$ of elements of A and a sequence $(S_i)_{i \in I}$ of elements of $\mathcal{S}$ such that one of the following holds:*

$$\text{(1)} \qquad \text{for all } i, j \in \omega \qquad a_i \in S_j \Leftrightarrow i < j$$

$$\text{(2)} \qquad \text{for all } i, j \in \omega \qquad a_i \in S_j \Leftrightarrow j < i.$$

PROOF. There are at most $|A|$ many pairs of finite subsets of A. So we may choose $\mathcal{S}' \subseteq \mathcal{S}$ with $|\mathcal{S}'| = |A|$ such that any two finite subsets A_0 and A_1 of A that can be separated by a set in $\mathcal{S}$ can also be separated by a set in $\mathcal{S}'$. By assumption there is some $S^* \in \mathcal{S}$ which is not a Boolean combination of elements of $\mathcal{S}'$, because there are at most $|B|$ different Boolean combinations of elements of $\mathcal{S}'$.

We construct by induction a sequence $(a'_i)_{i\in\omega}$ of elements of S^*, a sequence $(a''_i)_{i\in\omega}$ of elements of $B \setminus S^*$, and a sequence $(S_i)_{i\in\omega}$ of elements of $\mathcal{S}'$ such that for all $n \in \omega$

- $\{a'_0, \dots, a'_n\} \subseteq S_n$
- $\{a''_0, \dots, a''_n\} \subseteq A \setminus S_n$,
- $a'_n \in S_i \Leftrightarrow a''_n \in S_i$ for all $i < n$.

Assume that $(a'_i)_{i<n}$, $(a''_i)_{i<n}$, and $(S_i)_{i<n}$ have already been constructed. Since S^* is not a (positive) Boolean combination of $S_0, \dots, S_{n-1}$, there are $a'_n \in S^*$ and $a''_n \in A \setminus S^*$ such that for all $i < n$

$$a'_n \in S_i \Leftrightarrow a''_n \in S_i.$$

By the choice of $\mathcal{S}'$ there is some $S_n \in \mathcal{F}'$ with

$$\{a'_0, \dots, a'_n\} \subseteq S_n \text{ and } \{a''_0, \dots, a''_n\} \subseteq A \setminus S_n.$$

By Ramsey's theorem (Theorem C.0.2) we may assume that either

- $a'_j \notin S_i$ for all $i < j \in \omega$, or
- $a'_j \in S_i$ for all $i < j \in \omega$.

In the first case we set $(a_i)_{i\in\omega} := (a'_i)_{i\in\omega}$ and are in case (1): for all $i < j \in \omega$ we have $a'_i \in S_j$ and $a'_j \notin S_i$. In the second case, we set $(a_i)_{i\in\omega} := (a''_i)_{i\in\omega}$ and are in case (2): for all $i < j \in \omega$ we have $a''_i \notin S_j$ and $a''_j \in S_i$ since $a'_j \in S_i$. □

Bibliography

[1] S. A. Adeleke and P. M. Neumann. *Relations related to betweenness: their structure and automorphisms*, volume 623 of *Memoirs of the AMS*. American Mathematical Society, 1998.

[2] G. Ahlbrandt and M. Ziegler. Quasi-finitely axiomatizable totally categorical theories. *Annals of Pure and Applied Logic*, 30(1):63–82, 1986.

[3] M. Bodirsky. Introduction to mathematical logic, 2023. Course notes, TU Dresden, `https://wwwpub.zih.tu-dresden.de/~bodirsky/Logic.pdf`.

[4] M. Bodirsky, P. Jonsson, and T. V. Pham. The reducts of the homogeneous binary branching C-relation. *Journal of Symbolic Logic*, 81(4):1255–1297, 2016. Preprint arXiv:1408.2554.

[5] M. Bodirsky, M. Mamino, and C. Viola. Piecewise linear valued CSPs solvable by linear programming relaxation. *ACM Transactions of Computational Logic*, 23(1):1–35, 2022. Preprint arXiv:1912.09298.

[6] S. Braunfeld, A. Dawar, I. Eleftheriadis, and A. Papadopoulos. Monadic NIP in monotone classes of relational structures. In *50th International Colloquium on Automata, Languages, and Programming, ICALP 2023, July 10-14, 2023, Paderborn, Germany*, pages 119:1–119:18, 2023.

[7] P. J. Cameron. *Oligomorphic permutation groups*. Cambridge University Press, Cambridge, 1990.

[8] P. J. Cameron. Homogeneous permutations. *Electronic Journal of Combinatorics*, 9(2), 2002.

[9] A. Chernikov. *Lecture notes on stability theory*. AMS Open Math Notes, 2019.

[10] Dummit and Foote. *Abstract Algebra*. Wiley, 2004. Third edition.

[11] J. Ferrante and C. Rackoff. A decision procedure for the first order theory of real addition with order. *SIAM Journal on Computing*, 4(1):69–76, 1975.

[12] R. Fraïssé. Sur l'extension aux relations de quelques propriétés des ordres. *Annales Scientifiques de l'École Normale Supérieure*, 71:363–388, 1954.

[13] R. Fraïssé. *Theory of Relations*. Elsevier Science Ltd, North-Holland, Amsterdam, 1986.

[14] R. L. Graham, B. L. Rothschild, and J. H. Spencer. *Ramsey theory*. Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley & Sons, Inc., New York, 1990. Second edition.

[15] D. Haskell and D. Macpherson. Cell decompositions of C-minimal structures. *Annals of Pure and Applied Logic*, 66:113–162, 1994.

[16] C. W. Henson. Countable homogeneous relational systems and categorical theories. *Journal of Symbolic Logic*, 37:494–500, 1972.

[17] C. W. Henson. Model theory. Class Notes for Mathematics 571, University of Illinois, 2010.

[18] W. Hodges. *Model theory*. Cambridge University Press, Cambridge, 1993.

[19] W. Hodges. *A shorter model theory*. Cambridge University Press, Cambridge, 1997.

[20] T. Jech. *Set theory*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2003. The third millennium edition, revised and expanded.

[21] S. Lang. *Algebra*. Springer, 2002. Revised third edition.

[22] D. Marker. *Model Theory: An Introduction*. Springer, New York, 2002.

[23] Y. V. Matiyasevich. *Hilbert's Tenth Problem*. MIT Press, Cambridge, Massachusetts, 1993.

[24] B. Poonen. Characterizing integers among rational numbers with a universal-existential formula. *American Journal of Mathematics*, pages 675–682, 2009.

[25] P. Rothmaler. *Introduction to Model Theory*. CRC Press, 2000.

[26] D. Saracino. Model companions for $\aleph_0$-categorical theories. *Proceedings of the American Mathematical Society*, 39:591–598, 1973.

[27] J. Schreier and Stanisław Marcin Ulam. Über die Permutationsgruppe der natürlichen Zahlenfolge. *Studia Mathematica*, 4:134–141, 1933.

[28] P. Simon. *A Guide to NIP Theories*. Lecture Notes in Logic. Cambridge University Press, 2015.

[29] K. Tent and M. Ziegler. *A course in model theory*. Lecture Notes in Logic. Cambridge University Press, 2012.

[30] V. N. Vapnik and Z. Y. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Thoeret. Probl and its Appl.*, 16(2):264–280, 1971.