Math 509 Lecture Notes
Fall 2025

Alex Kruckman

Last updated December 18, 2025

Contents

1 Syntax and semantics

Lecture 1: 9/8

Mathematical logic, broadly speaking, is the field of mathematics which is concerned with the language we use to describe and reason about mathematical objects. That is, it highlights the distinction between syntax (the language) and semantics (the objects themselves). This may sound philosophical, but mathematicians constantly describe objects of interest by syntactic presentations, and then reason about the object by manipulating the syntax, e.g. the description of the circle as the set of solutions of the equation $x^2 + y^2 = 1$ in the plane, the coordinatization of a manifold by explicit charts, the specification of a construction or function by an algorithm, the presentation of a group by generators and relations, or the description of a function as an infinite series. If we go the extra step to viewing the syntax as a mathematical object to study in its own right, we begin doing mathematical logic.

There are many kinds of syntax in mathematics (as evidenced by the examples just given), and there are many logical systems. One of the most important is first-order logic, which is the system we will study in this course. Why is first-order logic important? For practical reasons: it is fairly expressive, while also exhibiting a number of nice properties, like the compactness theorem and a sound and complete proof system. And for foundational reasons: ZFC set theory, expressed in first-order logic, is the standard foundation for mathematics (but we will *not* discuss any foundational issues in this class).

Like any logic, first-order logic has a proof theory and a model theory. Proof theory is focused on syntax: systems for formal reasoning in the logic, and their properties. On the other hand, model theory is focused on semantics. The primary interest is in the relationships between properties of a first-order theory (a set of axioms which are sentences of first-order logic), its class of models (the mathematical structures satisfying the axioms), and definability within these models (elements, functions, sets, and relations definable by formulas of first-order logic).

Model theory is very abstract, in that we typically study first-order theories and structures in general, rather than any particular theory or structure. But we will often ground ourselves and obtain applications by specializing the general theory to particular examples.

Here, "mathematical structure" has a precise meaning: a structure is a set (or a family of sets), equipped with distinguished elements, operations, and relations. For example, we could consider the real numbers $\mathbb{R}$ as a structure in the language of order $\{<\}$, the language of abelian groups $\{0, +, -\}$, the language of fields $\{0, 1, +, -, \times\}$, or even larger languages suitable for elementary calculus, like $\{0, 1, +, -, \times, <, e^x, \sin(x)\}$. The choice of a language determines what we can say about the structure. For example, it determines a first-order theory with a class of models to study alongside $\mathbb{R}$: the class of dense linear orders, or torsion-free divisible abelian groups, or real closed fields. It also determines what is definable inside $\mathbb{R}$: rather than studying *all* functions $\mathbb{R} \to \mathbb{R}$ (or all continuous functions or all analytic functions), in model theory we study those functions which are definable in the language, i.e., those we can actually write down.

What is a first-order formula? Let me give some examples. Here is a formula in the language of groups:

$$x \cdot y = y \cdot x$$

For elements x and y In a group G, this formula expresses that x and y commute. Here is another:

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

In a group G, this formula expresses that G is abelian. In the first formula, x and y are free variables (we will say it is a formula in context $\{x, y\}$), so the formula expresses something about a pair of elements from a group. The second formula has no free variables (we will say it is a sentence), so the formula expresses something global about the group.

Here is a formula in the language of rings:

$$\exists y\, (x \cdot y = 1)$$

For an element x in a ring R, this formula expresses that x is a unit (i.e., x has a multiplicative inverse).

Here is a sentence in the language of orders:

$$\exists x\, \forall y\, (y = x \vee y < x)$$

In a linear order L, this formula expresses that the order has a greatest element.

How are the formulas above built? We can begin with variables, which refer to elements of a structure. From the variables and the distinguished elements and operations, we build terms which refer to other elements. The terms in the formulas above are x, y, $x \cdot y$, $y \cdot x$, and 1. Next, we make basic statements about the terms, called atomic formulas. The atomic formulas in the formulas above are $x \cdot y = y \cdot x$, $x \cdot y = 1$, $y = x$, and $y < x$. Finally, we put these atomic formulas together using logic connectives and quantifiers (which bind variables).

In first-order logic, the logical connectives are $\top$ and $\bot$ (for true and false), $\wedge$ (and), $\vee$ (lor), and $\neg$ (not). The quantifiers are $\forall$ (for all) and $\exists$ (exists).

Now, if we are given a structure (with explicit distinguished elements, operations, and relations) and an assignment of free variables to elements of the structure, we can make sense of any first-order formula in that structure in the obvious way.

What makes the logic "first-order" is that the variables and quantifiers range only over elements of the structure in question, not over higher-order objects like subsets or functions on that structure, and not over external sets like the natural numbers (unless the structure in question happens to be $\mathbb{N}$, of course).

At the beginning of the course, we need to make all this precise, defining the syntax and semantics of first-order logic one piece at a time, as shown in the table below.

Syntax	**Semantics**
Vocabularies	Structures
Variables	Assignments
Terms	Evaluation
Formulas	Satisfaction
Theories	Models

As the course proceeds, we will develop the basics of model theory assuming no prior background: the compactness and Löwenheim-Skolem theorems, realizing and omitting types, quantifier elimination, and the structure of countable models of complete theories. However, I will put an emphasis on some topics that are typically treated briefly, if at all, in a first course: fragments of first-order logic (especially the Horn, positive existential, and existential fragments), existentially closed structures, and model complete theories. For example, I plan to present a little-known proof of the compactness theorem, due to Poizat, that uses these concepts.

1.1 Languages and structures

Definition 1.1. A first-order **language** $\mathcal{L}$ is a set, whose elements are called **symbols**. Each symbol $s \in \mathcal{L}$ is designated as a **function symbol** or a **relation symbol**, and comes with an **arity** $\mathrm{ar}(s) \in \mathbb{N}$.

We are intentionally vague about what counts as a symbol; the name symbol is meant to suggest something that you could write down with a pencil on paper, but we have no intention of formalizing this notion. In practice, real-world symbols on paper can be encoded as mathematical objects (e.g. sets) in any way you like, and a symbol can be any mathematical object. In particular, a vocabulary may be uncountably infinite.

Intuitively, the arity of a symbol is the number of inputs it takes. The words **unary**, **binary**, and **ternary** mean arity 1, 2, and 3, respectively, and we also write n-**ary** to mean arity n. A 0-ary function symbol is called a **constant symbol**, and a 0-ary relation symbol is called a **proposition symbol**.

Example 1.2. The language of groups is $\mathcal{L}_{\text{Group}} = \{e, \cdot, ^{-1}\}$, where e is a constant symbol, $\cdot$ is a binary function symbol, and $^{-1}$ is a unary function symbol (for the inverse operation).

The language of rings is $\mathcal{L}_{\text{Ring}} = \{0, 1, +, -, \cdot\}$, where 0 and 1 are constant symbols, $+$ and $\cdot$ are binary function symbols, and $-$ is a unary function symbol (for the additive inverse operation).

The language of strict orders is $\mathcal{L}_{\text{SOrd}} = \{<\}$, where $<$ is a binary relation symbol. The language of orders is $\mathcal{L}_{\text{Ord}} = \{\leq\}$, where $\leq$ is a binary relation symbol. Note that the difference between $\mathcal{L}_{\text{SOrd}}$ and $\mathcal{L}_{\text{Ord}}$ is purely cosmetic.

The language of ordered rings is $\mathcal{L}_{\text{OrdRing}} = \mathcal{L}_{\text{Ring}} \cup \mathcal{L}_{\text{SOrd}} = \{0, 1, +, -, \cdot, <\}$.

For a ring R, the language of R-modules is $\mathcal{L}_{R\text{-Mod}} = \{0, +, -\} \cup \{\lambda_c \mid c \in R\}$, where 0 is a constant symbol, $+$ is a binary function symbol, $-$ is a unary function symbol, and each λ_c is a unary function symbol (for multiplication by the scalar c). This language is uncountable when R is, e.g., in the case of vector spaces over the real or complex numbers.

In all further discussions, we always have in the background a language $\mathcal{L}$, and we often omit mention of the language unless there a possibility for confusion, e.g., if there are multiple languages in play.

Definition 1.3. An **$\mathcal{L}$-structure** is a set A equipped with:

- For each function symbol $f \in \mathcal{L}$, a function $f^A \colon A^{\text{ar}(f)} \to A$.
- For each relation symbol $R \in \mathcal{L}$, a relation $R^A \subseteq A^{\text{ar}(f)}$.

In the case of a constant symbol c, we have $c^A \colon A^0 \to A$. The set A^0 is a singleton set $\{*\}$ (whose element $*$ is the unique 0-tuple from A, i.e., the empty sequence or empty function), and we identify c^A with the element $c^A(*) \in A$.

In the case of a proposition symbol P, we have $P^A \subseteq A^0 = \{*\}$, and P^A is either $\{*\}$ ("true") or $\varnothing$ ("false").

Contrary to a common convention, we do not require structures to be non-empty in general. But note that if $\mathcal{L}$ contains any constant symbols, then any $\mathcal{L}$-structure will be non-empty, since it must contain elements interpreting the constant symbols.

Example 1.4. As in Example 1.2, the language of rings is $\mathcal{L}_{\text{Ring}} = \{0, 1, +, -, \cdot\}$. Any ring R is an $\mathcal{L}_{\text{Ring}}$-structure in an obvious way, where 0^R is the zero element of the ring, $+^R$ is the addition operation in the ring, etc. Note, however, that not every $\mathcal{L}_{\text{Ring}}$-structure is a ring. Indeed, in an $\mathcal{L}_{\text{Ring}}$-structure, the symbols can be interpreted as arbitrary functions and relations. We will need to impose axioms to restrict ourselves to natural classes of structures.

It is also worth noting at this point that the interpretations of function symbols must be total. It is tempting to extend the language of rings to a language of fields by adding a unary function symbol $^{-1}$ for multiplicative inverse. But the interpretation of $^{-1}$ in a field K would have to be a total function $K \to K$, and 0 has no multiplicative inverse. One could make the ad hoc choice that $0^{-1} = 0$, but this can lead to unintuitive results later on interpretations of terms. For this reason, we typically use the language of rings when discussing fields; we will see later that the multiplicative inverse function is a definable function in this context.

1.2 Variables and assignments

Whenever we consider a term or formula, we would like to specify what free variables are in play. This is called the variable context of the term or formula. Since terms and formulas are finite syntactic objects, we are primarily interested in finite variable contexts, but it is sometimes useful to consider infinite contexts, so we do not restrict the definition to the finite case.

Definition 1.5. A **variable context** is a set, whose elements are called **variables**.

Just as with symbols, we will be intentionally vague about what counts as a variable, but we will follow the usual convention of writing them with letters near the end of the alphabet, like $x, y, z, \ldots$ or $x_0, x_1, x_2, \ldots$. We will often write a single letter, like x, for a variable context containing multiple variables. For example, we might have $x = \{x_0, \ldots, x_{n-1}\}$.

Definition 1.6. Given a variable context x and a structure A, an **assignment** of x in A is a function $a\colon x \to A$. We denote by A^x the set of all assignments of x in A.

If $x = \{x_0, \ldots, x_{n-1}\}$ is a variable context with n variables, and $a \in A^x$, we write a_i for $a(x_i)$, the element of A assigned to the variable x_i. There is a natural bijection between A^x and A^n, the set of n-tuples from A, carrying $a \in A^x$ to $(a_0, \ldots, a_{n-1})$ the n-tuple of interpretations of the variables $(x_0, \ldots, x_{n-1})$.

1.3 Terms and evaluation

Terms are syntactic expressions built from variables and function symbols that name elements of a structure.

Definition 1.7. An **$\mathcal{L}$-term** in context x is one of the following:

- A variable from x.
- A composite term $f(t_1, \ldots, t_n)$, where $f \in \mathcal{L}$ is an n-ary function symbol and $t_1, \ldots, t_n$ are $\mathcal{L}$-terms in context x.

We write $\mathrm{Term}_x(\mathcal{L})$ for the set of all $\mathcal{L}$-terms in context x, dropping $\mathcal{L}$ from the notation when it is clear from context.

Note that if $c \in \mathcal{L}$ is a constant symbol (i.e., a 0-ary function symbol), then c is a term. Indeed, it is the composite term formed from the function symbol c and 0 additional terms in context x.

The set Term_x is defined recursively, so we obtain methods of proof by induction and construction by recursion, with variables as the base case and the formation of composite terms as the inductive step. Sometimes it is useful to handle the constant symbols as a separate base case.

Example 1.8. In the language $\mathcal{L}_{\text{OrdRing}}$ of ordered rings, the following are terms in context $\{x, y\}$:

$$y, \quad 0, \quad (x+0)\cdot(-y), \quad ((x\cdot x)\cdot x)\cdot x, \quad (-(1+1))\cdot x$$

Of course, these are also terms in the language $\mathcal{L}_{\text{Ring}}$ of rings. The presence of the relation symbol $<$ is irrelevant for terms.

Note that we use the natural notation for our symbols when they differ from the formal syntax described above, for example writing $(x+0)$ instead of $+(x, 0)$. We use parentheses freely to avoid ambiguity. We will often to abbreviate terms in natural ways, e.g. by writing $((x\cdot x)\cdot x)\cdot x$ as x^4. But for this to make sense, we need to be working in a context where associativity of $\cdot$ is assumed.

Definition 1.9. Let A be an $\mathcal{L}$-structure, let t be an $\mathcal{L}$-term in context x, and let $a \in A^x$. Then we define the **evaluation** of t at a, written $t^A(a)$, by recursion on t:

- If t is a variable x_i in x, $t^A(a) = a(x_i) = a_i$.
- If t is a composite term $f(t_1, \dots, t_n)$, then we have elements $t_i^A(a) \in A$ by recursion for all $1 \leq i \leq n$. We define $t^A(a) = f^A(t_1^A(a), \dots, t_n^A(a))$.

Given a term $t \in \text{Term}_x$ and a structure A, there is a function $t^A \colon A^x \to A$, called a **term function**, defined by $a \mapsto t^A(a)$.

We often write $t(x)$ to denote a term t in context x (suggesting a function to which we can "plug in" an assignment $a \in A^x$ for the variables x). A term always comes with an associated variable context, even if it is not explicit in the notation.

If t is a term in context x, and $x \subseteq y$, then we can also view t as a term t' in the larger context y. Indeed, the context only restricts which variables can be mentioned. The term functions t^A and $(t')^A$ have different domains: A^x and A^y respectively. But they have the same behavior, in the sense that for all $b \in A^y$,

$$t^A(b|_x) = (t')^A(b).$$

Rigorously proving the assertions in the previous paragraph is a good straightforward exercise to make sure you understand proof by induction on terms.

1.4 Formulas and satisfaction

Lecture 2: 9/10

Formulas are syntactic expressions built from terms using $=$, relation symbols, and logical connectives and quantifiers. While terms evaluate to elements of a structure, formulas evaluate to "true" or "false".

Definition 1.10. An **atomic $\mathcal{L}$-formula** in context x is one of the following:

- $(t_1 = t_2)$, where t_1 and t_2 are $\mathcal{L}$-terms in context x.

- $R(t_1, \ldots, t_n)$, where $R \in \mathcal{L}$ is an n-ary relation symbol and $t_1, \ldots, t_n$ are $\mathcal{L}$-terms in context x.

We write $\mathrm{At}_x(\mathcal{L})$ for the set of all atomic $\mathcal{L}$-formula in context x, dropping $\mathcal{L}$ from the notation when it is clear from context.

Definition 1.11. An **$\mathcal{L}$-formula** in context x is one of the following:

- An atomic $\mathcal{L}$-formula in context x.
- $\top$ or $\bot$.
- $(\psi \wedge \chi)$, $(\psi \vee \chi)$, or $\neg\psi$, where ψ and χ are $\mathcal{L}$-formulas in context x.
- $\exists y\, \psi$ or $\forall y\, \psi$, where y is a variable not in x and ψ is an $\mathcal{L}$-formula in context $x \cup \{y\}$.

We write $\mathrm{Form}_x(\mathcal{L})$ for the set of all $\mathcal{L}$-formula in context x, dropping $\mathcal{L}$ from the notation when it is clear from context.

This is a recursive definition (simultaneously across all contexts x), so we obtain methods of proof by induction and construction by recursion, with atomic formulas and $\top$ and $\bot$ as base cases and the formation of formulas by boolean connectives and quantifiers as inductive steps.

As consequence of the way quantifiers and contexts interact in our definition, rebinding of variables is not allowed (contrary to the convention in some presentations of first-order logic). For example, $\exists y\, \forall y\, (x < y)$ is not a formula in context $\{x\}$, since $\forall y\, (x < y)$ is not a formula in context $\{x, y\}$.

Example 1.12. In the language $\mathcal{L}_{\mathrm{OrdRing}}$ of ordered rings, the following are formulas:

$$\text{Context } \{x, y, z\}\colon\ ((x \cdot x) + (x + x)) + 1 = 0, \quad x + z \leq y, \quad \exists w\, (x \cdot w = 1)$$
$$\text{Empty context } \varnothing\colon\ \neg(0 = 1), \quad \forall x\, (x = 0 \vee \exists y\, (x \cdot y = 1))$$

Note that we use the natural notation for our symbols when they differ from the formal syntax described above, for example writing $x \leq y$ instead of $\leq(x, y)$. In the context of rings, we could rewrite the first example as $x^2 + 2x + 1 = 0$.

We will also employ the following standard shorthands:

- $(t_1 \neq t_2)$ is shorthand for $\neg(t_1 = t_2)$.
- $(\psi \to \chi)$ is shorthand for $(\neg\psi \vee \chi)$.
- $(\psi \leftrightarrow \chi)$ is shorthand for $((\psi \to \chi) \wedge (\chi \to \psi))$.
- $\bigwedge_{i=1}^n \varphi_i$ and $\bigvee_{i=1}^n \varphi_i$ are shorthands for $(\ldots((\varphi_1 \wedge \varphi_2) \wedge \varphi_3) \cdots \wedge \varphi_n)$ and $(\ldots((\varphi_1 \vee \varphi_2) \vee \varphi_3) \cdots \vee \varphi_n)$, respectively. In the case $n = 0$, the empty conjunction is $\top$ and the empty disjunction is $\bot$.

- When $y = \{y_1, \ldots, y_n\}$ is a finite set of variables, $\exists y\, \varphi$ is shorthand for $\exists y_1 \ldots \exists y_n \varphi$, and $\forall y\, \varphi$ is shorthand for $\forall y_1 \ldots \forall y_n \varphi$. Technically we need to choose an order in which to quantify the variables, but this does not change the meaning of the formula.

The following definition, known as "Tarski's Definition of Truth" explains what it means for a structure to satisfy a formula by giving the various logical constructs their natural meanings (e.g., $\wedge$ means "and"). The result may seem trivial, but the point is that this is a formal definition by recursion and hence can be used to prove things by induction on formulas.

Definition 1.13. Let A be an $\mathcal{L}$-structure, let φ be an $\mathcal{L}$-formula in context x, and let $a \in A^x$. We define the relation $A \models \varphi(\overline{a})$, read A **satisfies** $\varphi(a)$ or φ is **true** of a in A, by recursion on φ:

- If φ is $(t_1 = t_2)$, then $A \models \varphi(a)$ iff $t_1^A(a) = t_2^A(a)$.
- If φ is $R(t_1, \ldots, t_n)$, then $A \models \varphi(a)$ iff $(t_1^A(a), \ldots, t_n^A(a)) \in R^A$.
- If φ is $\top$, then $A \models \varphi(a)$.
- If φ is $\bot$, then $A \not\models \varphi(a)$.
- If φ is $(\psi \wedge \chi)$, then $A \models \varphi(a)$ iff $A \models \psi(a)$ and $A \models \chi(a)$.
- If φ is $(\psi \vee \chi)$, then $A \models \varphi(a)$ iff $A \models \psi(a)$ or $A \models \chi(a)$.
- If φ is $\neg\psi$, then $A \models \varphi(a)$ iff $A \not\models \psi(a)$.
- If φ is $\exists y\, \psi$, then $A \models \varphi(a)$ iff there exists $b \in A$ such that $A \models \psi(a, b)$.
- If φ is $\forall y\, \psi$, then $A \models \varphi(a)$ iff for all $b \in A$, $A \models \psi(a, b)$.

In the quantifier clauses, ψ is a formula in context $x' = x \cup \{y\}$, and $A \models \psi(a, b)$ is shorthand for $A \models \psi(a')$, where $a' \in A^{x'}$ extends a by assigning the new variable y to b.

Given a formula $\varphi \in \mathrm{Form}_x$ and a structure A, we define

$$\varphi(A) = \{a \in A^x \mid A \models \varphi(a)\}$$

and call this a **definable set**.

We often write $\varphi(x)$ to denote a formula φ in context x. A formula always comes with an associated context, even if it is not explicit in the notation.

If φ is a formula in context x, and $x \subseteq y$ (with no variable in y appearing bound by a quantifier in φ), then we can also view φ as a formula φ' in the larger context y. The definable sets $\varphi(A)$ and $\varphi'(A)$ live in different Cartesian powers of A: $\varphi(A) \subseteq A^x$, while $\varphi'(A) \subseteq A^y$. But for all $b \in A^y$, we have

$$A \models \varphi(b|_x) \text{ if and only if } A \models \varphi'(b).$$

Again, proving these assertions is a good exercise in proof by induction on formulas.

1.5 Theories and models

Definition 1.14. An **$\mathcal{L}$-sentence** is an $\mathcal{L}$-formula in the empty context.

Some variables may appear in a sentence, but they must all be bound by quantifiers. When A is a structure, there is a unique assignment $* \in A^{\varnothing}$, so the satisfaction of a sentence φ does not depend on a choice of variable assignment. We write $A \models \varphi$ or $A \not\models \varphi$, instead of $A \models \varphi(*)$ or $A \not\models \varphi(*)$. Intuitively, a sentence expresses a property of A, rather than a property of tuples from A. In terms of definable sets, a sentence defines a subset of $A^{\varnothing}$, which is either $\{*\}$ ("true") or $\varnothing$ ("false").

Definition 1.15. An **$\mathcal{L}$-theory** is a set of $\mathcal{L}$-sentences. An $\mathcal{L}$-structure M is a **model** of an $\mathcal{L}$-theory T, written $M \models T$, if $M \models \varphi$ for all $\varphi \in T$.

We now overload the symbol $\models$ further.

Definition 1.16. If T is an $\mathcal{L}$-theory and φ is an $\mathcal{L}$-sentence, then T **entails** φ, written $T \models \varphi$, if every model of T satisfies φ.

Example 1.17. The $\mathcal{L}_{\text{Group}}$ theory of groups, T_{Group}, consists of the following three sentences:

$$\forall x \forall y \forall z\, ((x \cdot y) \cdot z = x \cdot (y \cdot z))$$
$$\forall x\, ((x \cdot e = x) \wedge (e \cdot x = x))$$
$$\forall x\, ((x \cdot x^{-1} = e) \wedge (x^{-1} \cdot x = e))$$

Of course, an $\mathcal{L}_{\text{Group}}$-structure G is a group if and only if $G \models T_{\text{Group}}$.

We have
$$T_{\text{Group}} \models \forall x\, (x \cdot x = e) \to \forall x \forall y\, (x \cdot y = y \cdot x),$$
since all groups of exponent 2 are abelian.

On the other hand, we have

$$T_{\text{Group}} \not\models \forall x \forall y\, (x \cdot y = y \cdot x),$$

since there exists a group which is not abelian.

2 Maps between structures

Lecture 3: 9/15

Now that we have set up a very general framework for studying mathematical structures, we want to establish some basic results which are common to structures in all languages, e.g., regarding substructures and homomorphisms between structures. In this section, we will also introduce the notion of a fragment of first-order logic, morphisms preserving the formulas in a fragment, and the method of diagrams.

2.1 Substructures

Suppose A is a subset of an $\mathcal{L}$-structure B. We say that A is **$\mathcal{L}$-closed** if it is closed under the functions f^A for all function symbols $f \in \mathcal{L}$, i.e., if $\operatorname{ar}(f) = n$, then for all $(a_1, \ldots, a_n) \in A^n$, we have $f^B(a_1, \ldots, a_n) \in A$.

If c is a constant symbol, the case $n = 0$ of the definition implies that $c^B \in A$.

If $\mathcal{L}$ is a relational language (i.e., it contains no function symbols), then every subset of B is $\mathcal{L}$-closed.

If $A \subseteq B$ is $\mathcal{L}$-closed, then we can turn A into an $\mathcal{L}$-structure, called the **induced substructure** on A, in a natural way:

- $f^A(a_1, \ldots, a_n) = f^B(a_1, \ldots, a_n)$ for each n-ary function symbol $f \in \mathcal{L}$ and each $(a_1, \ldots, a_n) \in A^n$.

- $(a_1, \ldots, a_n) \in R^A$ if and only if $(a_1, \ldots, a_n) \in R^B$ for each n-ary relation symbol $R \in \mathcal{L}$ and each $(a_1, \ldots, a_n) \in A^n$.

If general, if A and B are $\mathcal{L}$-structures and $A \subseteq B$, we say that A is a **substructure** of B if the interpretations of the symbols in $\mathcal{L}$ in A are those induced from B. That is, if A is $\mathcal{L}$-closed in B, $f^A = f^B|_{A^n}$ for each n-ary function symbol $f \in \mathcal{L}$, and $R^A = R^B \cap A^n$ for each n-ary relation symbol $R \in \mathcal{L}$.

Lemma 2.1. *Let B be an $\mathcal{L}$-structure, and let $A \subseteq B$ be an arbitrary subset. Then there is a smallest substructure of B containing A, denoted $\langle A \rangle_B$ and called the **substructure generated by** A. The underlying set of $\langle A \rangle_B$ is*

$$\{t^B(a) \mid t \textit{ is a term in context } x, \textit{ and } a \in A^x\}.$$

Proof. First, we show that $\langle A \rangle_B = \{t^B(a) \mid t \text{ is a term in context } x, \text{ and } a \in A^x\}$ is $\mathcal{L}$-closed. Let $f \in \mathcal{L}$ be an n-ary function symbol. Let $t_1^B(a_1), \ldots, t_n^B(a_n)$ be elements of $\langle A \rangle_B$. Each t_i is a term in context x_i, and $a_i \in A^{x_i}$. By renaming variables, we may assume that the x_i are pairwise disjoint. Let x be the context $\bigcup_{i=1}^n x_i$, and let $a \in A^x$ be the assignment which restricts to a_i on x_i. Then we can view each term t_i as a term t_i' in context x, and we have $(t_i')^B(a) = t_i^B(a_i)$. Let t be the composite term $f(t_1', \ldots, t_n')$ in context x. Now:

$$f^B(t_1^B(a_1), \ldots, t_n^B(a_n)) = f^B((t_1')^B(a), \ldots, (t_n')^B(a)) = t^B(a) \in \langle A \rangle_B,$$

so $\langle A\rangle_B$ is $\mathcal{L}$-closed. It follows that $\langle A\rangle_B$ is the underlying set of a substructure of B.

To see that $A \subseteq \langle A\rangle_B$, for each $a \in A$, let t be the term x in context $\{x\}$, and denote by a the assignment $x \mapsto a$ in A^x. Then $a = t^B(a) \in \langle A\rangle_B$, so $A \subseteq \langle A\rangle_B$.

Finally, we show that $\langle A\rangle_B$ is the smallest substructure of B containing A. Suppose C is a substructure of B containing A. To show $\langle A\rangle_B \subseteq C$, it suffices to show that $t^B(a) \in C$ for every term t in context x and $a \in A^x$. We prove this by induction on t.

If t is a variable x_i, then $t^B(a) = a_i \in A \subseteq C$.

If t is a composite term $f(t_1, \ldots, t_n)$, then $t^B(a) = f^B(t_1^B(a), \ldots, t_n^B(a))$. By induction, $t_i^B(a) \in C$ for all $1 \leq i \leq n$, and hence $f^B(t_1^B(a), \ldots, t_n^B(a)) \in C$, since C is $\mathcal{L}$-closed. □

We drop the subscript B from $\langle A\rangle_B$ when it is clear from context.

2.2 Homomorphisms and embeddings

Definition 2.2. If A and B are $\mathcal{L}$-structures, an **$\mathcal{L}$-homomorphism** $h\colon A \to B$ is a function such that:

- For every n-ary function symbol $f \in \mathcal{L}$ and for every tuple $(a_1, \ldots, a_n)$ from A, $h(f^A(a_1, \ldots, a_n)) = f^B(h(a_1), \ldots, h(a_n))$.
- For every n-ary relation symbol $R \in \mathcal{L}$ and for every tuple $(a_1, \ldots, a_n)$ from A, if $(a_1, \ldots, a_n) \in R^A$, then $(h(a_1), \ldots, h(a_n)) \in R^B$. We say h **preserves** R.

If $c \in \mathcal{L}$ is a constant symbol, the case $n = 0$ in the definition implies that $h(c^A) = c^B$.

Definition 2.3. A homomorphism $h\colon A \to B$ is an **$\mathcal{L}$-embedding** if:

- h is injective.
- For every n-ary relation symbol $R \in \mathcal{L}$ and for every tuple $(a_1, \ldots, a_n)$ from A, if $(h(a_1), \ldots, h(a_n)) \in R^B$, then $(a_1, \ldots, a_n) \in R^A$. We say h **reflects** R.

For most classes mathematical structures (groups, rings, graphs, etc.), the general notion of homomorphism defined above agrees with the specialized notion of homomorphism defined for that class (group homomorphisms, ring homomorphisms, graph homomorphisms, etc.). But it is sensitive to the choice of language.

For example, if we consider the language $\mathcal{L} = \{0, +, -, \cdot\}$ (without a constant symbol for 1), there are $\mathcal{L}$-homomorphisms $R \to S$ between unital rings which fail to map 1^R to 1^S.

For another example, any homomorphism between two linear orders in the language $\mathcal{L}_{\mathrm{SOrd}} = \{<\}$ must be strictly increasing (i.e., $a < b$ implies $h(a) <$

$h(b)$) and hence injective. But a homomorphism between two linear orders in the language $\mathcal{L}_{\text{Ord}} = \{\leq\}$ need only be weakly increasing (i.e., $a < b$ implies $h(a) \leq h(b)$) and may fail to be injective.

Embeddings are largely important because of their connection to the notion of substructure: If A and B are substructures and $A \subseteq B$, then A is a substructure of B if and only if the inclusion map $i\colon A \to B$ is an embedding.

Homomorphisms and embeddings are defined in terms of the symbols in the language. A natural question is: how do they interact with the rest of our logical structure? First, we show they commute with the evaluation of terms.

Given a function $h\colon A \to B$ and an assignment $a \in A^x$, we write $h(a)$ for $(h \circ a) \in B^x$. If $x = \{x_0, \ldots, x_{n-1}\}$ and a assigns the variables in x to the elements $(a_0, \ldots, a_{n-1})$, then $h(a)$ assigns the variables in x to the elements $(h(a_0), \ldots, h(a_{n-1}))$.

Lemma 2.4. *Let $h\colon A \to B$ be a homomorphism. Then for any term $t(x)$ and any assignment $a \in A^x$, we have*

$$h(t^A(a)) = t^B(h(a)).$$

Proof. By induction on t.

If t is a variable x_i in x, then $h(t^A(a)) = h(a_i) = t^B(h(a))$.

If t is a composite term $f(t_1, \ldots, t_n)$, then

$$\begin{aligned} h(t^A(a)) &= h(f^A(t_1^A(a), \ldots, t_n^A(a))) \\ &= f^B(h(t_1^A(a)), \ldots, h(t_n^A(a))) && h \text{ is a homomorphism} \\ &= f^B(t_1^B(h(a)), \ldots, t_n^B(h(a))) && \text{by induction} \\ &= t^B(h(a)). && \square \end{aligned}$$

Definition 2.5. Let A and B be $\mathcal{L}$-structures, and let $h\colon A \dashrightarrow B$ be a partial function, i.e., a function $C \to B$ for some set $C \subseteq A$. We say that h **preserves** an $\mathcal{L}$-formula $\varphi(x)$ when for any $c \in C^x$, if $A \models \varphi(c)$, then $B \models \varphi(h(c))$. We say that h **reflects** $\varphi(x)$ when for any $c \in C^x$, if $B \models \varphi(h(c))$, then $A \models \varphi(c)$.

Remark 2.6. A partial function h preserves φ if and only if it reflects $\neg\varphi$.

Indeed, h reflects $\neg\varphi$ if and only if $B \models \neg\varphi(h(a))$ implies $A \models \neg\varphi(a)$ if and only if $A \models \varphi(a)$ implies $B \models \varphi(h(a))$ if and only if h preserves φ.

You can think of the next proposition as a generalization of the conditions for defining a homomorphism of groups $G \to H$ by defining a function on the generators of G. To be well-defined, the images of the generators in H have to satisfy all the same relations as the generators do in G.

Proposition 2.7. *Let A and B be $\mathcal{L}$-structures, and let $h\colon A \dashrightarrow B$ be a partial function defined on a set of generators for A. That is, the domain of h is $C \subseteq A$ with $A = \langle C \rangle_A$. Then h extends to a homomorphism $h'\colon A \to B$ if and only if h preserves all atomic formulas. Moreover, in this case h' is unique.*

Proof. Suppose h extends to a homomorphism $h'\colon A \to B$. Let $\varphi(x)$ be an atomic formula, and let $c \in C^x$. We show that h preserves φ.

Case 1: φ is $t_1 = t_2$. If $A \models \varphi(c)$, then $t_1^A(c) = t_2^A(c)$, so $h'(t_1^A(c)) = h'(t_2^A(c))$. By Lemma 2.4, $h'(t_i^A(c)) = t_i^B(h'(c)) = t_i^B(h(c))$ for $i \in \{1,2\}$, so $t_1^B(h(c)) = t_2^B(h(c))$, and $B \models \varphi(h(c))$.

Case 2: φ is $R(t_1,\dots,t_n)$. If $A \models \varphi(c)$, then $(t_1^A(c),\dots,t_n^A(c)) \in R^A$, so $(h'(t_1^A(c)),\dots,h'(t_n^A(c))) \in R^B$, since h' preserves R. Again, by Lemma 2.4, $h'(t_i^A(c)) = t_i^B(h'(c)) = t_i^B(h(c))$ for $1 \le i \le n$, so $(t_1^B(h(c)),\dots,t_n^B(h(c))) \in R^B$, and $B \models \varphi(h(c))$.

Conversely, suppose h preserves all atomic formulas. We will show that h extends to a homomorphism $h'\colon A \to B$. Since $A = \langle C\rangle_A$, by Lemma 2.1, every element of A can be written as $t^A(c)$ for some term t in context x and some $c \in C^x$. Note that given finitely many elements of A, we can assume that they have the form $t_1^A(c),\dots,t_n^A(c)$, where each term t_i is in the same context x and $c \in C^x$, by expanding the context of each term as in the proof of Lemma 2.1.

We define $h'(t^A(c)) = t^B(h(c))$. I claim this is well-defined. If $t_1^A(c) = t_2^A(c)$, then $A \models (t_1 = t_2)(c)$, which implies $B \models (t_1 = t_2)(h(c))$, since h preserves atomic formulas. So $t_1^B(h(c)) = t_2^B(h(c))$, and h' is well-defined. Lecture 4: 9/17

Next, we check that h' commutes with the function symbols in the language. Let $f \in \mathcal{L}$ be an n-ary function symbol, and let $(a_1,\dots,a_n) \in A^n$ with $a_i = t_i^A(c)$ for each i. Let s be the composite term $f(t_1,\dots,t_n)$. Then:

$$\begin{aligned}
h'(f^A(a_1,\dots,a_n)) &= h'(f^A(t_1^A(c),\dots,t_n^A(c)))\\
&= h'(s^A(c))\\
&= s^B(h(c))\\
&= f^B(t_1^B(h(c)),\dots,t_n^B(h(c))\\
&= f^B(h'(t_1^A(c)),\dots,h'(t_n^A(c)))\\
&= f^B(h'(a_1),\dots,h'(a_n)).
\end{aligned}$$

Finally, let $R \in \mathcal{L}$ be an n-ary relation symbol, and let $(a_1,\dots,a_n) \in A^n$ with each $a_i = t_i^A(c)$. Let $\varphi(x)$ be the atomic formula $R(t_1,\dots,t_n)$. Then $(a_1,\dots,a_n) \in R^A$ if and only if $(t_1^A(c),\dots,t_n^A(c)) \in R^A$ if and only if $A \models \varphi(c)$. Since h preserves atomic formulas, this implies $B \models \varphi(h(c))$, which means $(t_1^B(h(c)),\dots,t_n^B(h(c))) = (h'(t_1^A(c)),\dots,h'(t_n^A(c))) = (h'(a_1),\dots,h'(a_n)) \in R^B$. Thus h' is a homomorphism.

It remains to show that h' is unique. Suppose $h''\colon A \to B$ is a homomorphism extending h. Let $a \in A$, and suppose $a = t^A(c)$ for some $c \in C^x$. By Lemma 2.4, $h''(a) = h''(t^A(c)) = t^A(h''(c)) = t^A(h(c)) = h'(t^A(c)) = h'(a)$, so $h'' = h'$. □

It is an immediate consequence of Proposition 2.7 that a total function $h\colon A \to B$ is a homomorphism if and only if it preserves all atomic formulas. We now prove the corresponding characterization of embeddings. A **literal** is an atomic or negated atomic formula.

Proposition 2.8. *Let A and B be $\mathcal{L}$-structures and $h\colon A \to B$ a function. The following are equivalent:*

(1) h is an embedding.

(2) h preserves and reflects all atomic formulas.

(3) h preserves all literals.

Proof. The equivalence of (2) and (3) follows from Remark 2.6: Reflecting atomic formulas is equivalent to preserving negated atomic formulas.

Assume h preserves and reflects all atomic formulas. Since h preserves all atomic formulas, it is a homomorphism by Lemma 2.7. Since h reflects atomic formulas of the form $x = y$, it is injective. Since h reflects atomic formulas of the form $R(x_1, \dots, x_n)$ with R and n-ary relation symbol, it reflects relation symbols.

Now suppose h is an embedding. Since h is a homomorphism, it preserves atomic formulas. Let $\varphi(x)$ be an atomic formula, let $a \in A^x$, and assume $B \models \varphi(h(a))$. If φ is $t_1 = t_2$, then $t_1^B(h(a)) = t_2^B(h(a))$, so by Lemma 2.4, $h(t_1^A(a)) = h(t_2^A(a))$. Since h is injective, $t_1^A(a) = t_2^A(a)$, so $A \models \varphi(a)$. If φ is $R(t_1, \dots, t_n)$, then $(t_1^B(h(a)), \dots, t_n^B(h(a))) \in R^B$. By Lemma 2.4, $(t_1^B(h(a)), \dots, t_n^B(h(a))) = (h(t_1^A(a)), \dots, h(t_n^A(a)))$, and since h reflects relation symbols, $(t_1^A(a), \dots, t_n^A(a)) \in R^A$, so $A \models \varphi(a)$. $\square$

2.3 Fragments and $\mathcal{F}$-morphisms

Given a map between structures, once can ask the question of which formulas it preserves and reflects. Conversely, given a set of formulas, one can focus on maps preserving these formulas.

Definition 2.9. A **fragment** of first-order logic is a set $\mathcal{F}$ of $\mathcal{L}$-formulas which contains all atomic $\mathcal{L}$-formulas and is closed under subformula and substitution of terms for free variables.

Usually, we consider fragments that contain all atomic formulas and are built recursively from these using specified formula-building operations. For fragments of this form, the closure conditions will hold trivially.

Here are some examples of fragments:

- The **atomic fragment** (At) contains all atomic formulas.
- The **literal fragment** (Lit) contains all literals (atomic and negated atomic formulas).
- The **positive quantifier-free fragment** (qf^+) contains formulas built from At using $\wedge$ and $\vee$.
- The **quantifier-free fragment** (qf) contains formulas built from At using $\wedge$, $\vee$, and $\neg$.

- The **positive primitive fragment** (pp) contains formulas built from At using $\wedge$ and $\exists$. In categorical logic, this is called the **regular** fragment.
- The **positive existential fragment** ($\exists^+$) contains formulas built from At using $\wedge$, $\vee$, and $\exists$. In categorical logic, this is called the **coherent** fragment.
- The **existential fragment** ($\exists$) contains formulas built from Lit using $\wedge$, $\vee$, and $\exists$.
- The **universal fragment** ($\forall$) contains formulas built from Lit using $\wedge$, $\vee$, and $\forall$.
- The **elementary fragment** (FO) is the set of all first-order formulas.

Given a fragment $\mathcal{F}$, let $\exists(\mathcal{F})$ be the closure of $\mathcal{F}$ under the formula-building operations of $\wedge$, $\vee$, and $\exists$. So, for example, the positive existential fragment is $\exists(\text{At})$, while the existential fragment is $\exists(\text{Lit})$. Similarly, let $\forall(\mathcal{F})$ be the closure of $\mathcal{F}$ under the formula-building operations of $\wedge$, $\vee$, and $\forall$. We then define a heirarchy of fragments stratifying FO by quantifier complexity:

$$\begin{aligned}
\exists_1 &= \exists(\text{Lit}) \\
\forall_1 &= \forall(\text{Lit}) \\
\exists_{n+1} &= \exists(\forall_n) \text{ for all } n \\
\forall_{n+1} &= \forall(\exists_n) \text{ for all } n.
\end{aligned}$$

Our convention above is that when we close a fragment under $\wedge$, we include $\top$ (the empty conjunction), and when we close a fragment under $\vee$, we include $\bot$ (the empty disjunction).

Often, we only care about membership in a fragment up to logical equivalence.

Definition 2.10. Two $\mathcal{L}$-formulas $\varphi(x)$ and $\psi(x)$ in the same context are **logically equivalent** if $A \models \varphi(a)$ if and only if $A \models \psi(a)$ for all $\mathcal{L}$-structures A and all $a \in A^x$.

It is a classical result from propositional logic that every quantifier-free formula is logically equivalent to one in disjunctive normal form: $\bigvee_{i=1}^{n} \bigwedge_{j=1}^{m} \varphi_{i,j}$, where $\varphi_{i,j}$ is a literal for all i and j. It follows that, up to equivalence, we could have described the quantifier-free fragment as the closure of Lit under $\wedge$ and $\vee$. It also follows that, up to equivalence, we could have described the existential fragment as the closure of qf (rather than Lit) under $\wedge$, $\vee$, and $\exists$, and similarly for all the $\exists_n$ and $\forall_n$.

Existential formulas are frequently defined have the form $\exists y_1 \dots \exists y_n \varphi(x, y)$, where φ is quantifier-free, i.e., the existential fragment is the closure of qf under $\exists$. The justification for this is prenex normal form: every formula is equivalent to one where all quantifiers are pulled to the front. But prenex normal form is not valid over empty structures, so we have to use a different convention. For example, when P is a proposition symbol, $(\exists x\,\top) \vee P$ is an existential formula, but it is not logically equivalent to any formula in prenex normal form. Indeed, $(\exists x\,\top)\vee P$ is true in an empty structure when P is true, but no formula beginning with an existential quantifier is true in an empty structure.

Definition 2.11. Given a fragment $\mathcal{F}$, an **$\mathcal{F}$-morphism** is a function which preserves all formulas in $\mathcal{F}$.

Note that for any fragment $\mathcal{F}$, since we assume $\mathcal{F}$ contains all atomic formulas, every $\mathcal{F}$-morphism is a homomorphism by Proposition 2.7.

By Remark 2.6, an $\mathcal{F}$-morphism reflects all negations of formulas in $\mathcal{F}$. If $\mathcal{F}$ is closed under $\neg$ (up to equivalence), then a $\mathcal{F}$-morphism both preserves and reflects all formulas in $\mathcal{F}$.

It follows from Proposition 2.7 that the At-morphisms are exactly the homomorphism, and it follows from Proposition 2.8 that the Lit-morphisms are exactly the embeddings.

We call an FO-morphism an **elementary embedding**. At this point, it is not clear that any non-trivial elementary embeddings exist! However, there is one kind of morphism that preserves *all* structure.

Definition 2.12. An **isomorphism** is a homomorphism $h\colon A \to B$ such that there exists an inverse homomorphism $h^{-1}\colon B \to A$. We write $A \cong B$ when A and B are isomorphic, i.e., when there exists an isomorphism $A \to B$.

It is a good exercise to prove the following:

(a) A homomorphism $h\colon A \to B$ is an isomorphism if and only if it is a surjective embedding.

(b) Every isomorphism is an elementary embedding.

It turns out that if a function preserves $\mathcal{F}$-formulas, then it preserves some others for free.

Theorem 2.13. *Let $h\colon A \to B$ be an $\mathcal{F}$-morphism. Then h is an $\exists(\mathcal{F})$-morphism.*

Proof. Let $\varphi(x)$ be a formula in $\exists(\mathcal{F})$. We prove that h preserves φ by induction on φ. Let $a \in A^x$, and assume $A \models \varphi(a)$.

In the base case, $\varphi \in \mathcal{F}$. Then h preserves φ by hypothesis.

If φ is $\top$ or $\bot$, then h preserves φ trivially.

Suppose φ is $\psi\wedge\chi$. Then $A \models \psi(a)$ and $A \models \chi(a)$. By induction, h preserves ψ and χ, so $B \models \psi(h(a))$ and $B \models \chi(h(a))$, and thus $B \models \varphi(h(a))$.

Suppose φ is $\psi \vee \chi$. Then $A \models \psi(a)$ or $A \models \chi(a)$. By induction, h preserves ψ and χ, so $B \models \psi(h(a))$ or $B \models \chi(h(a))$, and thus $B \models \varphi(h(a))$.

Suppose φ is $\exists y\,\psi(x, y)$. Then there is some $b \in A$ such that $A \models \psi(a, b)$. By induction, h preserves ψ, so $B \models \psi(h(a), h(b))$. Thus $B \models \exists y\,\psi(h(a), y)$, i.e., $B \models \varphi(h(a))$. $\square$

As a consequence of the theorem, all homomorphisms preserve the positive existential fragment $\exists^+$ and all embeddings preserve the existential fragment $\exists$.

Remark 2.14. By De Morgan's Laws, the negation of a universal formula is logically equivalent to an existential formula and vice versa. Since embeddings preserve all existential formulas, embeddings reflect all negated existential formulas, and hence reflect all universal formulas.

The quantifier-free fragment is contained (up to equivalence) in both the existential and universal fragments, so embeddings both preserve and reflect all quantifier-free formulas.

Example 2.15. Consider the structure $(\mathbb{Z}; <)$ and the map $h\colon \mathbb{Z} \to \mathbb{Z}$ given by $n \mapsto 2n$. This is an embedding: it is injective, and $a < b$ if and only if $2a < 2b$.

Let $\varphi(x, y)$ be the formula $\exists z\,(x < z < y)$. ($x < y < z$ is shorthand for $x < z \wedge z < y$.) This is an existential formula. It is preserved by h, since if there is some c such that $a < c < b$, then there is some c' such that $2a < c' < 2b$. For example, we can take $c' = 2c$. On the other hand, φ is not reflected by h. For example, $\mathbb{Z} \models \varphi(h(1), h(2))$, since $2 < 3 < 4$, but $\mathbb{Z} \not\models \varphi(1, 2)$.

Example 2.16. In $\mathcal{L}_{\text{Group}} = \{\cdot, e, ^{-1}\}$, the group axioms are all given by universal sentences (e.g., $\forall x\,(x \cdot e = x \wedge e \cdot x = x)$). It follows that if G is a group and $H \subseteq G$ is a substructure, then H is a group, since the axioms are reflected by the inclusion embedding.

If we tried to axiomatize groups in the language $\mathcal{L} = \{\cdot\}$, we would have to use axioms with existential quantifiers that are not reflected to substructures (e.g., $\exists z \forall x\,(x \cdot z = x \wedge z \cdot x = x)$). And indeed, there are $\mathcal{L}$-substructures of groups which are not groups, like $(\mathbb{N}; +) \subseteq (\mathbb{Z}; +)$.

Let G be a group, and let H be a subgroup of G. The formula $x \cdot y = y \cdot x$ expressing that two elements commute is preserved and reflected by the inclusion embedding: two elements commute in H if and only if they commute in G. The formula $\forall y\,(x \cdot y = y \cdot x)$ defining the center is reflected by the inclusion embedding, but not preserved: $Z(H)$ need not be contained in $Z(G)$, but $Z(G) \cap H \subseteq Z(H)$. Finally, the universal sentence $\forall x\,\forall y\,(x \cdot y = y \cdot x)$ is reflected by the inclusion embedding: this just says that a subgroup of an abelian group is abelian.

2.4 Diagrams

Given an $\mathcal{L}$-structure A and a subset $C \subseteq A$, let $\mathcal{L}(C)$ be the language obtained from $\mathcal{L}$ by adding a new constant symbol for every element $c \in C$. When there is no chance for confusion, we will also denote the constant symbol by c. We view A as an $\mathcal{L}(C)$ structure in the obvious way, with $c^A = c$. Often we do not distinguish notationally between A as a $\mathcal{L}$-structure and as a $\mathcal{L}(C)$-structure, but when we need to we will denote the $\mathcal{L}(C)$-structure by A_C.

When $\mathcal{L} \subseteq \mathcal{L}'$ are languages and A is an $\mathcal{L}'$-structure, we write $A|_{\mathcal{L}}$ for the **reduct** of A to the language $\mathcal{L}$, obtained by forgetting about the interpretations of the symbols in $\mathcal{L} \setminus \mathcal{L}'$.

Example 2.17. For any $C \subseteq A$, $A_C|_{\mathcal{L}} = A$.

Example 2.18. Let $\mathcal{L}_{\text{Ab}} = \{0, +, -\}$ be the language of abelian groups, and let $\mathcal{L}_{\text{Ring}} = \{0, 1, +, -, \cdot\}$ be the language of rings. If R is a ring, the reduct $R|_{\mathcal{L}}$ is the underlying abelian group of R.

Suppose C is a subset of a structure A. Given a formula $\varphi(x)$ and an assignment $c \in C^x$, we will denote by $\varphi(c)$ the $\mathcal{L}(C)$-sentence obtained by substituting for each variable x_i appearing in $\varphi(x)$ the constant symbol c_i corresponding to the element if C assigned to x_i.

For any fragment $\mathcal{F}$, the the **$\mathcal{F}$-diagram** of C in A, denoted $\text{Diag}_A^{\mathcal{F}}(C)$, is the set of all $\mathcal{L}(C)$-sentences $\varphi(c)$ such that $\varphi(x) \in \mathcal{F}$ and $A \models \varphi(c)$.

When $\mathcal{F}$ is the atomic fragment, $\text{Diag}_A^{\text{At}}(C)$ is called the **positive diagram** and denoted $\text{Diag}_A^{+}(C)$. When $\mathcal{F}$ is the literal fragment, $\text{Diag}_A^{\text{Lit}}(C)$ is just called the **diagram** and denoted $\text{Diag}_A(C)$. When $\mathcal{F}$ is the elementary fragment, $\text{Diag}_A^{\text{FO}}(C)$ is called the **elementary diagram**. We drop the subscript A when it is clear from context.

Proposition 2.19. *Let A be an $\mathcal{L}$-structure, and let B be a $\mathcal{L}(A)$-structure. Then $B \models \text{Diag}^{\mathcal{F}}(A)$ if and only if the function $h\colon A \to B$ given by $a \mapsto a^B$ is an $\mathcal{F}$-morphism $A \to B|_{\mathcal{L}}$.*

Proof. For any formula $\varphi(x)$ in $\mathcal{F}$, we have that h preserves $\varphi(x)$ if and only if for all $a \in A^x$, if $A \models \varphi(a)$, then $B \models \varphi(a)$. Equivalently, if $\varphi(a) \in \text{Diag}^{\mathcal{F}}(A)$, then $B \models \varphi(a)$.

Now h is an $\mathcal{F}$-morphism if and only if it preserves all formulas in $\mathcal{F}$. As we have just shown, this is equivalent to the condition that for all $\varphi(a) \in \text{Diag}^{\mathcal{F}}(A)$, $B \models \varphi(a)$, i.e., $B \models \text{Diag}^{\mathcal{F}}(A)$. □

The significance of Proposition 2.19 is that we can turn the problem of finding some $\mathcal{F}$-morphism from a structure A to another structure with certain properties into the problem of finding a model for some $\mathcal{L}(A)$ theory, namely an appropriate diagram. We can build homomorphisms using $\text{Diag}^{+}(A)$, embeddings using $\text{Diag}(A)$, and elementary embeddings using $\text{Diag}_{\text{FO}}(A)$.

3 Horn theories and initial models

3.1 Free, initial, and terminal structures

Lecture 5: 9/22

Many examples of $\mathcal{L}$-structures come from mathematical practice. If we are interested in groups, or rings, or modules, we may decide on an appropriate language $\mathcal{L}$, axiomatize them by an $\mathcal{L}$-theory T, and consider these objects as models of T.

But if we are interested in an *arbitrary* theory T, where do we get models of T? How do we know there are any at all? In fact there may not be, since T might be inconsistent. But an insight of mathematical logic is that we can often build models of T directly from the syntax of T.

As a first example of what I mean, I will show how to build $\mathcal{L}$-structures (without worrying about satisfying any theory T) from terms.

Let x be a variable context. Recall that we denote by Term_x the set of all $\mathcal{L}$-terms in context x. We make Term_x into an $\mathcal{L}$-structure, called the **term algebra** in context x, as follows:

- For each n-ary function symbol $f \in \mathcal{L}$ and $t_1, \ldots, t_n \in \mathrm{Term}_x$,
$$f^{\mathrm{Term}_x}(t_1, \ldots, t_n) = f(t_1, \ldots, t_n) \in \mathrm{Term}_x.$$
- For each relation symbol $R \in \mathcal{L}$, $R^{\mathrm{Term}_x} = \varnothing$.

There is a function $x \to \mathrm{Term}_x$, mapping each variable in x to its corresponding term. Viewing this function as a variable assignment, the variables in x are assigned to themselves, so we also denote this assignment by x. It is easy to prove by induction that for any term $t \in \mathrm{Term}_x$, $t^{\mathrm{Term}_x}(x) = t$.

Now let $\varphi(x)$ be an atomic formula. If φ is $t_1 = t_2$, then $\mathrm{Term}_x \models \varphi(x)$ if and only if $t_1^{\mathrm{Term}_x}(x) = t_2^{\mathrm{Term}_x}(x)$ if and only if $t_1 = t_2$. If φ is $R(t_1, \ldots, t_n)$, then $\mathrm{Term}_x \not\models \varphi(x)$ (since $R^{\mathrm{Term}_x} = \varnothing$).

The motivation for interpreting every relation symbol as the empty relation is to make it as easy as possible to define homomorphisms out of the term algebra. We want Term_x to satisfy as few atomic formulas as possible.

Proposition 3.1. *The term algebra* Term_x *is the free* $\mathcal{L}$*-structure on the set* x*. That is, for every* $\mathcal{L}$*-structure* A *and every function* $a\colon x \to A$*, there is a unique homomorphism* $\mathrm{eval}_a\colon \mathrm{Term}_x \to A$ *such that* $\mathrm{eval}_a \circ x = a$*.*

Proof. It is straightforward to prove directly that $\mathrm{eval}_a(t) = t^A(a)$ is the desired homomorphisms, but it is slicker to apply Proposition 2.7. Since every element $t \in \mathrm{Term}_x$ is the evaluation $t^{\mathrm{Term}_x}(x)$, we have $\langle x\rangle = \mathrm{Term}_x$, i.e., x is a set of generators for Term_x.

Now $a\colon x \to A$ is a partial function $\mathrm{Term}_x \to A$. As we noted above, if $\varphi(x)$ is an atomic formula such that $\mathrm{Term}_x \models \varphi(x)$, then φ has the form $t = t$, and automatically $A \models \varphi(a)$. Thus a preserves atomic formulas and extends to a unique homomorphism $\mathrm{eval}_a\colon \mathrm{Term}_x \to A$, defined by $\mathrm{eval}_a(t) = t^A(a)$. □

Corollary 3.2. *For any language $\mathcal{L}$, the term algebra* $\mathrm{Term}_\varnothing$ *in the empty context is an initial $\mathcal{L}$-structure. That is, it admits a unique homomorphism to every other $\mathcal{L}$-structure.*

Proof. For any $\mathcal{L}$-structure A, there is a unique function $*\colon \varnothing \to A$, which extends to a unique homomorphism $\mathrm{Term}_\varnothing \to A$. □

A term in the empty context (i.e., containing no variables) is called a **closed term**. Note that if $\mathcal{L}$ has no constant symbols, then there are no closed terms, and $\mathrm{Term}_\varnothing$ is an empty structure.

Of course, if we have an $\mathcal{L}$-theory T in mind, the initial $\mathcal{L}$-structure $\mathrm{Term}_\varnothing$ will typically fail to be a model of T. For example, the initial $\mathcal{L}_{\mathrm{Group}}$-structure has distinct elements e, $e \cdot e$, $e^{-1} \cdot ((e \cdot e)^{-1} \cdot e)$, etc. If we want to construct an initial model of the theory of groups, we need to collapse these distinct terms to be a single element $\{e\}$. This is what we will do in the next subsection. First, though, let's note that there is also a terminal $\mathcal{L}$-structure.

Let $\mathbf{1}$ be the structure with a single element $\{*\}$, defined as follows:

- For each n-ary function symbol $f \in \mathcal{L}$, $f^{\mathbf{1}}(*, \dots, *) = *$.
- For each n-ary relation symbol $R \in \mathcal{L}$, $R^{\mathbf{1}} = \mathbf{1}^n = \{(*, \dots, *)\}$.

For each term $t(x)$, there is a unique assignment $*\colon x \to \mathbf{1}$, and we have $t^{\mathbf{1}}(*) = *$ (the term must evaluate to some element, and this is the only one).

For each atomic formula $\varphi(x)$, if φ is $t_1 = t_2$, then $t_1^{\mathbf{1}}(a) = * = t_1^{\mathbf{1}}(a)$, so $\mathbf{1} \models \varphi(a)$. And if φ is $R(t_1, \dots, t_n)$, then $\mathbf{1} \models \varphi(a)$, since $R^{\mathbf{1}} = \mathbf{1}^n$. So $\mathbf{1}$ satisfies every atomic formula in every variable assignment.

Proposition 3.3. *The trivial structure* $\mathbf{1}$ *is the terminal $\mathcal{L}$-structure. That is, for every $\mathcal{L}$-structure A, there is a unique homomorphism $A \to \mathbf{1}$.*

Proof. It is straightforward to prove directly that the unique function $*\colon A \to \mathbf{1}$ is a homomorphism. We can also apply Proposition 2.7: as we noted above, $\mathbf{1}$ satisfies every atomic formula in every variable assignment, so $*$ preserves all atomic formulas and therefore is a homomorphism. □

3.2 Horn theories

A **sequent** is an expression of the form $\varphi \vdash_x \psi$, where x is a finite variable context and φ and ψ are formulas in context x. Given an $\mathcal{L}$-structure A, we define $A \models (\varphi \vdash_x \psi)$ to mean that for all assignments $a \in A^x$, if $A \models \varphi(a)$, then $A \models \psi(a)$. Equivalently, $\varphi(A) \subseteq \psi(A)$.

In full first-order logic, the sequent $\varphi \vdash_x \psi$ has the same content as the sentence $\forall x\, (\varphi \to \psi)$ (where $\forall x$ is shorthand for the block of universal quantifiers quantifying over all the variables in x). But in a fragment $\mathcal{F}$ of first-order logic that lacks an equivalent of $\to$ (i.e., lacks $\neg$ or $\vee$) or $\forall$, the sequent may not be equivalent to any sentence. And we still may want to assert that one $\mathcal{F}$-definable set is always contained in another.

Given a fragment $\mathcal{F}$ of first-order logic, we define an $\mathcal{F}$-sequent to be an expression of the form $\varphi \vdash_x \psi$, where φ and ψ are both in $\mathcal{F}$. By default, an **$\mathcal{F}$-theory** is a set of $\mathcal{F}$-sequents.

We now focus on the $\wedge$ fragment of first-order logic, whose formulas are conjunctions of atomic formulas (including the empty conjunction $\top$). A $\wedge$-sequent has the form

$$\varphi_1 \wedge \cdots \wedge \varphi_m \vdash_x \psi_1 \wedge \cdots \wedge \psi_n$$

where each φ_i and ψ_i is an atomic formula. Note that the single sequent above has the same content as the following n sequents, taken together:

$$\begin{gathered} \varphi_1 \wedge \cdots \wedge \varphi_m \vdash_x \psi_1 \\ \vdots \\ \varphi_1 \wedge \cdots \wedge \varphi_m \vdash_x \psi_n \end{gathered}$$

(and if $n = 0$, the sequent $\varphi_1 \wedge \cdots \wedge \varphi_m \vdash_x \top$ has no content – it is trivially satisfied in every structure).

For this reason, it is traditional to restrict attention to $\wedge$-sequents in which the right-hand-side consists of a single atomic formula. Such a sequent is called a **Horn clause**. We will call a set of Horn clauses (i.e., a $\wedge$-theory) a **Horn theory**.

In many sources, what I call Horn clauses are called "strict Horn clauses". I will explain the reason for this later. You will often see Horn clauses presented in the form $\forall x\, ((\varphi_1 \wedge \cdots \wedge \varphi_n) \to \psi)$ where each φ_i and ψ is atomic, or, worse, $\forall x\, (\varphi_1 \vee \cdots \vee \varphi_m)$ where exactly one of the φ_i is atomic and the rest are negated atomic. I think these presentations make the notion seem ad hoc, obscuring the fact that Horn clauses are really the natural notion of an axiom in the $\wedge$-fragment of first-order logic.

In a Horn clause, the left-hand-side is allowed to be the empty conjunction $\top$. A Horn clause of the form $\top \vdash_x \psi$ has the same content as the universally quantified atomic formula $\forall x\, \psi$. We call it an **atomic axiom**, and in the case when ψ is an equation $t_1 = t_2$, we call it an **equational axiom**. The field of universal algebra is primarily concerned with equational theories, i.e., those axiomatized by equational axioms. We usually write atomic axioms as $\vdash_x \psi$, omitting the $\top$.

Example 3.4. In $\mathcal{L}_{\text{Group}}$, the theory of groups is an equational theory:

$$\begin{aligned} &\vdash_{\{x,y,z\}} (x \cdot y) \cdot z = x \cdot (y \cdot z) \\ &\vdash_{\{x\}} x \cdot e = x \\ &\vdash_{\{x\}} e \cdot x = x \\ &\vdash_{\{x\}} x \cdot x^{-1} = e \\ &\vdash_{\{x\}} x^{-1} \cdot x = e \end{aligned}$$

Similarly, the theories of rings and of R-modules are equational theories in the appropriate languages.

The theory of torsion-free groups (those with no non-identity elements of finite order) is not equational, but it is a Horn theory, obtained by adding to the equational theory of groups the following infinite schema of axioms, one for each $n > 1$:

$$x^n = e \vdash_{\{x\}} x = e$$

Similarly, the theory of reduced rings (those with no non-zero nilpotent elements) is not equational, but it is a Horn theory, obtained by adding to the equational theory of rings the following infinite schema of axioms, one for each $n > 1$:

$$x^n = 0 \vdash_{\{x\}} x = 0$$

In fact, it is sufficient to add the single axiom

$$x^2 = 0 \vdash_{\{x\}} x = 0.$$

Indeed, if R satisfies this latter axiom, we can prove by induction on $n > 1$ that if $x^n = 0$, then $x = 0$. The base case $n = 2$ is assumed. Suppose $n > 0$ and $x^n = 0$. Then $(x^{n-1})^2 = x^n(x^{n-2}) = 0$, so by our axiom $x^{n-1} = 0$, and by induction $x = 0$.

The theories of integral domains and of fields are not even Horn theories. However, they can be axiomatized in the positive quantifier-free fragment qf^+ and in the positive existential fragment $\exists^+$, respectively. The theory of integral domains extends the equational theory of commutative rings by two further axioms.

$$\begin{aligned} 0 = 1 &\vdash_\varnothing \bot \\ x \cdot y = 0 &\vdash_{\{x,y\}} x = 0 \vee y = 0 \end{aligned}$$

The theory of fields existends the equational theory of commutative rings by two further axioms.

$$\begin{aligned} 0 = 1 &\vdash_\varnothing \bot \\ \top &\vdash_{\{x\}} x = 0 \vee \exists y\, (x \cdot y = 1). \end{aligned}$$

The terminal $\mathcal{L}$-structure $\mathbf{1}$ is a model of every Horn $\mathcal{L}$-theory. Indeed, since $\mathbf{1}$ satisfies all atomic formulas in all variable assignments, every Horn clause $\bigwedge_{i=1}^n \varphi_i \vdash_x \psi$ is trivially satisfied in $\mathbf{1}$.

In light of Example 3.4, $\mathbf{1}$ is the trivial group and the zero ring. But the theories of integral domains and fields do not admit $\mathbf{1}$ as a model, since $\mathbf{1}$ fails to satisfy $0 = 1 \vdash_\varnothing \bot$. This proves the assertion that these theories do not admit Horn axiomatizations.

Example 3.5. In $\mathcal{L}_{\mathrm{SOrd}} = \{<\}$, the class of (strict) posets does not admit a Horn axiomatization. Indeed, the $\mathcal{L}_{\mathrm{SOrd}}$-structure $\mathbf{1}$ is not a strict poset, since $* < *$ in this structure.

However, in the language $\mathcal{L}_{\text{Ord}} = \{\leq\}$, we can axiomatize partial orders by a Horn theory:

$$\vdash_{\{x\}} x \leq x$$
$$x \leq y \wedge y \leq z \vdash_{\{x,y,z\}} x \leq z$$
$$x \leq y \wedge y \leq x \vdash_{\{x,y\}} x = y.$$

The class of linear orders cannot be axiomatized by a Horn theory, even in the language $\mathcal{L}_{\text{Ord}}$. We need an additional qf$^+$-axiom:

$$\vdash_{\{x,y\}} (x \leq y) \vee (y \leq x)$$

3.3 Initial models of Horn theories

Lecture 6: 9/24

The immediate relevance of Horn theories for us is that they are the broadest context in which terminal and initial models are guaranteed to exist. We have already seen the terminal structure $\mathbf{1}$ is a model of every Horn theory T, and therefore is the terminal model of T. The initial models are much more interesting.

Definition 3.6. Let Δ be a set of atomic sentences. We say Δ is **diagrammatic** if it satisfies the following closure conditions:

(1) For every term t, $t = t \in \Delta$.

(2) For all terms t and t', if $t = t' \in \Delta$, then $t' = t \in \Delta$.

(3) For all terms t_1, t_2, and t_3, if $t_1 = t_2 \in \Delta$ and $t_2 = t_3 \in \Delta$, then $t_1 = t_3 \in \Delta$.

(4) For every n-ary function symbol $f \in \mathcal{L}$ and all terms $t_1, \ldots, t_n$ and $t'_1, \ldots, t'_n$, if $t_i = t'_i \in \Delta$ for all $1 \leq i \leq n$, then $f(t_1, \ldots, t_n) = f(t'_1, \ldots, t'_n) \in \Delta$.

(5) For every n-ary relation symbol $R \in \mathcal{L}$ and all terms $t_1, \ldots, t_n$ and $t'_1, \ldots, t'_n$, if $t_i = t'_i \in \Delta$ for all $1 \leq i \leq n$ and $R(t_1, \ldots, t_n) \in \Delta$, then $R(t'_1, \ldots, t'_n) \in \Delta$.

Let T be a Horn theory. We say Δ is **T-diagrammatic** if it additionally satisfies:

(6) For every Horn clause $\bigwedge_{i=1}^{n} \varphi_i \vdash_x \psi$, and every assignment $t \in \text{Term}_{\varnothing}^{x}$ of a closed term to each variable in x, if $\varphi_i(t) \in \Delta$ for all $1 \leq i \leq n$, then $\psi(t) \in \Delta$. Here $\varphi_i(t)$ is the atomic sentence obtained by substituting for the variables in x the corresponding terms in t.

Lemma 3.7. *Suppose Δ is a diagrammatic set of atomic sentences. Then there exists a structure M_Δ with $\langle \varnothing \rangle = M$ such that $\Delta = \text{Diag}^{+}_{M_\Delta}(\varnothing)$. Moreover, if Δ is T-diagrammatic, then $M_\Delta \models T$.*

Proof. We define a relation $\sim_\Delta$ on $\text{Term}_{\varnothing}$ by $t_1 \sim_\Delta t_2$ if and only if $t_1 = t_2 \in \Delta$. By closure conditions (1), (2), and (3), $\sim_\Delta$ is an equivalence relation. We write $[t]$ for the equivalence class of term t. Let $M_\Delta = \text{Term}_{\varnothing}/{\sim}$.

We now turn M_Δ into an $\mathcal{L}$-structure as follows:

- For each n-ary function symbol $f \in \mathcal{L}$,

$$f^{M_\Delta}([t_1], \dots, [t_n]) = [f(t_1, \dots, t_n)].$$

- For each n-ary relation symbol $R \in \mathcal{L}$,

$$R^{M_\Delta} = \{([t_1], \dots, [t_n]) \in M_\Delta^n \mid R(t_1, \dots, t_n) \in \Delta\}.$$

By closure conditions (4) and (5), these functions and relations are well-defined.

It is easy to show by induction that for every closed term t, $t^{M_\Delta} = [t]$. Since every element of M_Δ is the evaluation of some closed term, $M_\Delta = \langle \varnothing \rangle$.

For an atomic sentence of the form $t_1 = t_2$, we have $t_1 = t_2 \in \Delta$ if and only if $[t_1] = [t_2]$ if and only if $t_1^{M_\Delta} = t_2^{M_\Delta}$ if and only if $t_1 = t_2 \in \mathrm{Diag}^+_{M_\Delta}(\varnothing)$.

For an atomic sentence of the form $R(t_1, \dots, t_n)$, we have $R(t_1, \dots, t_n) \in \Delta$ if and only if $([t_1], \dots, [t_n]) \in R^{M_\Delta}$ if and only if $R(t_1, \dots, t_n) \in \mathrm{Diag}^+_{M_\Delta}(\varnothing)$.

Thus $\Delta = \mathrm{Diag}^+_{M_\Delta}(\varnothing)$. It remains to show that $M_\Delta \models T$, assuming Δ is T-diagrammatic.

Let $\bigwedge_{i=1}^n \varphi_i \vdash_x \psi$ be a Horn clause in T. Assuming $x = \{x_1, \dots, x_m\}$, let $([t_1], \dots, [t_m]) \in M_\Delta^x$ and assume $M_\Delta \models \bigwedge_{i=1}^n \varphi_i([t_1], \dots, [t_m])$. Since for each $1 \leq i \leq n$, $M_\Delta \models \varphi_i([t_1], \dots, [t_m])$, the atomic sentence $\varphi_i(t_1, \dots, t_m)$ is true in M_Δ. Thus $\varphi_i(t_1, \dots, t_n) \in \Delta$. By the T-diagrammatic condition, $\psi(t_1, \dots, t_n) \in \Delta$, so $M_\Delta \models \psi([t_1], \dots, [t_n])$. We have shown $M_\Delta \models \bigwedge_{i=1}^n \varphi_i \vdash_x \psi$, so $M_\Delta \models T$. □

Lemma 3.8. *For any structure M, $\Delta = \mathrm{Diag}^+_M(\varnothing)$ is diagrammatic. Moreover, if $M \models T$, then Δ is T-diagrammatic.*

Proof. It is clear that Δ is diagrammatic. Assuming $M \models T$, we verify (6). Let $\bigwedge_{i=1}^n \varphi_i \vdash_x \psi$ be a Horn clause in T, $t \in \mathrm{Term}^x_\varnothing$, and $\varphi_i(t) \in \Delta$ for all $1 \leq i \leq n$. Then $M \models \varphi_i(t^M)$ for all $1 \leq i \leq n$, and therefore $M \models \bigwedge_{i=1}^n \varphi_i(t^M)$. Since $M \models T$, $M \models \psi(t^M)$, so $\psi(t) \in \Delta$. □

Lemma 3.9. *Let T be a Horn theory. Then there is a least T-diagrammatic set of atomic sentences, Δ_T. Moreover, $\Delta_T = \bigcup_{T_0 \subseteq_{\mathrm{fin}} T} \Delta_{T_0}$. That is, if $\delta \in \Delta_T$, then there is a finite subset $T_0 \subseteq T$ such that $\delta \in \Delta_{T_0}$.*

Proof. We just need to take the least set of atomic sentences closed under conditions (1)–(6). Precisely, let $\Delta_0 = \varnothing$, and given Δ_n, let Δ_{n+1} be Δ_n together with all atomic sentences required by conditions (1)–(6). For example (condition 2), if $t = t' \in \Delta_n$, then we put $t' = t$ in Δ_{n+1}. For another example (condition 6), if $\bigwedge_{i=1}^n \varphi_i \vdash_x \psi$ is a Horn clause in T and $\varphi_i(t) \in \Delta_n$ for all $1 \leq i \leq n$, then we put $\psi(t)$ in Δ_{n+1}. Then we let $\Delta_T = \bigcup_{n \in \mathbb{N}} \Delta_n$.

Since each of the closure conditions (1)–(6) have the form "if some finite number of atomic sentences is in Δ, then some other atomic sentence is in Δ", it is easy to see that Δ_T is T-diagrammatic. Indeed, the finitely many sentences will be in Δ_N for some sufficiently large N, and then the resulting sentence will be in Δ_{N+1}. If Δ' is another T-diagrammatic set of atomic sentences, then clearly $\Delta_0 \subseteq \Delta'$, and by induction $\Delta_n \subseteq \Delta'$ for all n, so $\Delta_T \subseteq \Delta'$.

Now suppose $\delta \in \Delta_T$. Then there is some least n such that $\delta \in \Delta_n$. We prove by induction on n that $\delta \in \Delta_{T_0}$ for some finite $T_0 \subseteq T$. Since $\delta \notin \Delta_{n-1}$, it must have been added to Δ_n because of one of the closure conditions (1)–(6). Then there are finitely many $\delta_1, \ldots, \delta_k \in \Delta_{n-1}$ which witnessed the addition of δ to Δ_n. For each δ_i, by induction there is some finite $T^i \subseteq T$ such that $\delta_i \in \Delta_{T^i}$. Let $T_0 = \bigcup_{i=1}^k T^i$, and if δ was added because of closure condition (6) due to some horn clause $C \in T$, also add C to T_0.

Now since each $\delta_i \in \Delta_{T^i}$, $\delta_1, \ldots, \delta_k \in \Delta_{T_0}$. Since Δ_{T_0} is T-diagrammatic (and contains clause C, if we are in case (6)), we must have $\delta \in \Delta_{T_0}$ as well. □

Theorem 3.10. *Let T be a Horn theory. Then T has an initial model M_T (i.e., $M_T \models T$ and if $N \models T$, then there is a unique homomorphism $M_T \to N$).*

Proof. Let Δ be the least T-diagrammatic set of atomic sentences by Lemma 3.9. Let $M_T = M_{\Delta_T}$ from Lemma 3.7. Then $M_T = \langle \varnothing \rangle$ and $\Delta_T = \mathrm{Diag}^+_{M_T}(\varnothing)$.

Let $N \models T$. By Lemma 3.8, $\mathrm{Diag}^+_N(\varnothing)$ is T-diagrammatic, so $\mathrm{Diag}^+_{M_T}(\varnothing) = \Delta_T \subseteq \mathrm{Diag}^+_N(\varnothing)$. It follows that the unique function $\varnothing \to N$ preserves atomic formulas and hence extends to a unique homomorphism $M \to N$ by Proposition 2.7. □

Corollary 3.11. *Let T be a Horn theory. For an atomic sentence δ, the following are equivalent:*

(1) $\delta \in \Delta_T$.

(2) $M_T \models \delta$.

(3) $T \models \delta$.

(4) There is some finite subset $T_0 \subseteq T$ such that $T_0 \models \delta$.

Proof. (1) $\Leftrightarrow$ (2): Since $M_T = M_{\Delta_T}$, we have $\mathrm{Diag}^+_{M_T}(\varnothing) = \Delta_T$. Thus $\delta \in \Delta_T$ if and only if $M_T \models \delta$.

(2) $\Rightarrow$ (3): Suppose $M_T \models \delta$ and $N \models T$ is arbitrary. Since M_T is the initial model of T, there is a unique homomorphism $h\colon M_T \to N$. Since h preserves atomic sentences, $N \models \delta$. Thus $T \models \delta$.

(3) $\Rightarrow$ (2): $T \models \delta$ means every model of T satisfies δ. Since $M_T \models T$, $M_T \models \delta$.

(3) $\Rightarrow$ (4): Suppose $T \models \delta$. By the equivalence of (1) and (3), $\delta \in \Delta_T$. By Lemma 3.9, there is some finite subset T_0 such that $\delta \in \Delta_{T_0}$. By the equivalence of (1) and (3) again, $T_0 \models \delta$.

(4) $\Rightarrow$ (3): If $T_0 \subseteq T$ and $T_0 \models \delta$, then for any M such that $M \models T$, also $M \models T_0$, so $M \models \delta$. □

The equivalence of (3) and (4) in Corollary 3.11 is our first manifestation of the compactness theorem. We will deduce the full compactness theorem for first-order logic by bootstrapping from this version for atomic sentences relative to Horn theories.

One can view our work above as providing a complete proof system for deriving atomic sentences from Horn theories. It has the following proof rules:

$$\frac{}{t = t} \qquad \frac{t_2 = t_1}{t_1 = t_2} \qquad \frac{t_1 = t_2, t_2 = t_3}{t_1 = t_3}$$

$$\frac{t_1 = t_1', \cdots, t_n = t_n'}{f(t_1, \ldots, t_n) = f(t_1', \ldots, t_n')} \; f \in \mathcal{L} \text{ an } n\text{-ary function symbol}$$

$$\frac{t_1 = t_1', \cdots, t_n = t_n', R(t_1, \ldots, t_n)}{R(t_1', \ldots, t_n')} \; R \in \mathcal{L} \text{ an } n\text{-ary relation symbol}$$

$$\frac{\varphi_1(t_1, \ldots, t_n), \cdots \varphi_m(t_1, \ldots, t_n)}{\psi(t_1, \ldots, t_n)} \; \bigwedge_{i=1}^m \varphi_i \vdash_x \psi \in T$$

3.4 Models presented by generators and relations

In algebra, it is common to present a structure by generators and relations. For example, the dihedral group D_n can be presented as

$$\langle \sigma, \tau \mid \sigma^n = e, \tau^2 = e, \tau\sigma = \sigma^{-1}\tau \rangle.$$

We can view this presentation as a set of constant symbols C, together with a set Δ of atomic $\mathcal{L}_{\text{Group}}(C)$-sentences.

The group G so presented can be defined by a universal property. For any other group H together with a map $h\colon C \to H$ such that the images of C in H satisfy the relations in Δ, there is a unique homomorphism $h'\colon G \to H$ extending h. Since we can view the map $h\colon C \to H$ as an expansion of the group H to an $\mathcal{L}_{\text{Group}}(C)$-structure, we can rephrase this in our model-theoretic language: For any $\mathcal{L}_{\text{Group}}(C)$-structure $H \models T_{\text{Group}} \cup \Delta$, there is a unique $\mathcal{L}_{\text{Group}}(C)$-homomorphism $G \to H$. That is, G is the initial model of the theory $T_{\text{Group}} \cup \Delta$.

It may seem awkward that $T_{\text{Group}} \cup \Delta$ consists of two kinds of things: Horn (in fact, equational) axioms and atomic sentences. But recall that any atomic sentence δ has the same content as the sequent $\vdash_\varnothing \delta$, which is an atomic axiom in the empty context.

Making this all more precise: Given a Horn theory T in the language $\mathcal{L}$, a set of new constant symbols C, and a set of atomic $\mathcal{L}(C)$-sentences Δ, let $T(\Delta)$ be the Horn theory

$$T \cup \{\vdash_\varnothing \delta \mid \delta \in \Delta\}.$$

The **model presented by generators** C **and relations** Δ is the initial model $M_{T(\Delta)}$, which exists by Theorem 3.10.

Note that the proof of Lemma 3.9 can be easily adapted to show that for any set Δ of atomic sentences, there is a smallest T-diagrammatic set containing Δ. But this is just the same as the minimal $T(\Delta)$-diagrammatic set. The additional closure conditions from the atomic axioms $\vdash_\varnothing \delta$ in $T(\Delta)$ amount to the same thing as requiring a $T(\Delta)$-diagrammatic set to contain Δ.

Example 3.12. Consider the Horn theory of posets in $\mathcal{L}_{\text{Ord}} = \{\leq\}$. Suppose we have a presentation by generators and relations (C, Δ). For simplicity, let's assume that $c = c' \notin \Delta$ when c and c' are distinct constants. Then Δ amounts to a binary relation R on C (with cRc' if and only if $c \leq c' \in \Delta$). Closing up to the diagram of $M_{T(\Delta)}$ is equivalent to taking the reflexive and transitive closure of this R to obtain a preorder $\preceq$, and then passing to the poset obtained as a quotient by the equivalence relation $\sim$ defined by $c \sim c'$ iff $c \preceq c'$ and $c' \preceq c$. Lecture 7: 9/29

What is the cardinality of the model $M_{T(\Delta)}$? We cannot hope to provide a precise formula: even in the case of the theory of groups, it is an undecidable problem whether the initial model of a given finite set of atomic sentences Δ is the trivial group. However, we can easily provide an (infinite) upper bound.

Remark 3.13. Note that for any language $\mathcal{L}$, a closed $\mathcal{L}$-term is a finite sequence of symbols from $\mathcal{L}$ (possibly together with finitely many other symbols like parentheses and commas). The cardinality of the set of such finite sequences is bounded above by $\max(\aleph_0, |\mathcal{L}|)$. Since every element of M_T is named by a closed term, $|M_T| \leq \max(\aleph_0, |\mathcal{L}|)$.

Since $T(\Delta)$ expands the language by a new set of constant symbols C, we have $|\mathcal{L}(C)| = |\mathcal{L}| + |C| \leq \max(\aleph_0, |\mathcal{L}|, C)$, so $|M_{T(\Delta)}| \leq \max(\aleph_0, |\mathcal{L}|, |C|)$. For example, a finitely generated group is always at most countably infinite, and a group with a generating set of size κ (an infinite cardinal) has cardinality at most κ.

3.5 Horn theories with constraints

We now extend our scope very slightly by adding the contradictory formula $\perp$ to the $\wedge$-fragment. This allows us to form sequents of the form $\perp \vdash_x \psi$ (which we ignore, since they are trivially satisfied by every structure) and $\bigwedge_{i=1}^n \varphi_i \vdash_x \perp$. We call a sequent of the latter form a **constraint clause**. A theory in this fragment is called a **Horn theory with constraints**. Other sources have the convention that constraint clauses are also Horn clauses, while Horn clauses without $\perp$ are called "strict Horn clauses".

Note that we have $M \models (\bigwedge_{i=1}^n \varphi_i \vdash_x \perp)$ if and only if there is no assignment $a \in M^x$ satisfying $\bigwedge_{i=1}^n \varphi_i$. So constraint clauses forbid certain configurations from occurring in their models.

Unlike Horn theories, which are guaranteed to have models, Horn theories with constraints may be inconsistent. For example, the contradictory constraint clause $\top \vdash_\varnothing \perp$ has no models. Less trivially, $T_{\text{Group}} \cup \{x \cdot x = e \vdash_{\{x\}} \perp\}$ has no models, since every group has an element which squares to the identity (namely e itself).

Theorem 3.14 (Compactness for Horn theories with constraints)**.** *Let T be a Horn theory with constraints. Suppose that whenever $T_0 \subseteq T$ contains finitely many Horn clauses and exactly one constraint clause, T_0 has a model. Then T has a model. In fact, the initial model of the set of Horn clauses in T is an initial model of T.*

Proof. Write $T = T_H \cup T_C$, where T_H consists of the Horn clauses in T and T_C consists of the constraint clauses in T. Let M be the initial model of T_H. Toward a contradiction, assume $M \not\models T$. Then there is some constraint clause $C \in T_C$ such that $M \not\models C$. Write C as $\bigwedge_{i=1}^n \varphi_i \vdash_x \bot$. Then there is some $a \in M^x$ such that $M \models (\bigwedge_{i=1}^n \varphi_i)(a)$. Since $M = \langle \varnothing \rangle$, there exists $t \in \mathrm{Term}_\varnothing^x$ such that $a = t^M$. Then for all $1 \leq i \leq n$, $M \models \varphi_i(t)$. By Corollary 3.11, for all $1 \leq i \leq n$, there is a finite subtheory $T^i \subseteq T_H$ such that $T^i \models \varphi_i(t)$.

Let $T_0 = \{C\} \cup \bigcup_{i=1}^n T^i$. Then T_0 contains finitely many Horn clauses and exactly one constraint clause. By our hypothesis, T_0 has a model $N \models T_0$. Since $N \models T^i$ for all $1 \leq i \leq n$, we have $N \models \varphi_i(t)$ for all $1 \leq i \leq n$, so $N \models (\bigwedge_{i=1}^n \varphi_i)(t^N)$, contradicting $N \models C$.

Thus M is a model of T. For any other $N \models T$, we have $N \models T_H$, so there is a unique homomorphism $M \to N$. Thus M is an initial model of T. □

As an immediate corollary, we get a version of the Löwenheim–Skolem theorem in this context. This corollary implies that for any class of algebraic structures which is Horn-axiomatizable in a countable language, any infinite set can be the domain of a structure in the class.

Corollary 3.15 (Löwenheim–Skolem for Horn theories)**.** *Let T be a Horn theory, and assume that T has a model M with $|M| \geq 2$. Then for every infinite cardinal $\kappa \geq |\mathcal{L}|$, T has a model of cardinality κ.*

Proof. Given an infinite cardinal κ, let C be a set of new constant symbols of cardinality κ. Let T' be the following Horn theory with constraints in the language $\mathcal{L}(C)$:

$$T \cup \{c = c' \vdash_\varnothing \bot \mid c \neq c' \text{ in } C\}.$$

Suppose $T_0 \subseteq T'$ is a subset containing finitely many Horn clauses and exactly one constraint clause. This constraint clause is $c \neq c'$ for some distinct c and c' in C. We can expand M to a model of T_0 by interpreting c and c' as distinct elements of M and interpreting the rest of the constant symbols in C arbitrarily. Let M_T be the initial model of T when viewed as an $\mathcal{L}(C)$-theory. By Remark 3.13, $|M_T| \leq \max(\aleph_0, |\mathcal{L}|, |C|) = \kappa$. By Theorem 3.14, $M_T \models T'$, so the interpretations of the constants in C are distinct in M_T, and $|M_T| \geq \kappa$. Thus $|M_T| = \kappa$, and clearly $M_T|_{\mathcal{L}} \models T$, as desired. □

4 Existential theories and e.c. models

4.1 Direct limits

In the previous section, we constructed initial models for Horn theories. One can think of these models as "minimalist" in the sense that the only things that happen in them (in the sense of the elements that exist and the atomic formulas they satisfy) are what is forced to happen by the theory. We will now turn to constructing existentially closed models, which are "maximalist": in a sense, everything that could happen does happen in these models.

We will build these models as direct limits. The direct limit construction (which in category theory is called the directed colimit) involves pasting together a family of structured indexed by a directed set.

Definition 4.1. A **directed set** is a non-empty poset $(I, \leq)$ with the property that for all $i, j \in I$, there exists $k \in I$ with $i \leq k$ and $j \leq k$.

For example, any non-empty linear order is a directed set (we can take $k = \max(i, j)$). Another example is the set of all finite subsets of a non-empty set, ordered by $\subseteq$ (we can take $k = i \cup j$).

Definition 4.2. Let $(I, \leq)$ be a directed set. A **directed family** indexed by I is a family of $\mathcal{L}$-structures $(M_i)_{i\in I}$ together with a family of homomorphisms $f_{ij} \colon M_i \to M_j$ for all $i \leq j$ in I, such that

- For all $i \in I$, $f_{ii} = \mathrm{id}_{M_i}$.
- For all $i \leq j \leq k \in I$, $f_{jk} \circ f_{ij} = f_{ik}$.

In categorical language, a directed family is a functor from the poset $(I, \leq)$ (viewed as a category) to the category of $\mathcal{L}$-structures and homomorphisms, and the direct limit is the colimit of this diagram. It will be useful to have a concrete description of the elements of the direct limit, so we give an explicit construction rather than a category-theoretic characterization.

Given a directed family $((M_i)_{i\in i}, (f_{ij})_{i\leq j})$, we define the direct limit $\varinjlim M_i$. First, let $M_* = \bigsqcup_{i\in I} M_i$, the disjoint union of all structures in the family. We define a relation $\sim$ on M_*: given $a \in M_i$ and $b \in M_j$, $a \sim b$ if and only if there exists $k \in I$ with $i \leq k$ and $j \leq k$ such that $f_{ik}(a) = f_{jk}(b)$. It is clear that $\sim$ is reflexive and symmetric, so we check transitivity.

Suppose $a_1 \in M_{i_1}$, $a_2 \in M_{i_2}$, and $a_3 \in M_{i_3}$ with $a_1 \sim a_2$ and $a_2 \sim a_3$. Then there exist $j_1, j_2 \in I$ with $i_1 \leq j_1$, $i_2 \leq j_1$, $i_2 \leq j_2$, and $i_3 \leq j_2$ such that $f_{i_1j_1}(a_1) = f_{i_2j_1}(a_2)$ and $f_{i_2j_2}(a_2) = f_{i_3j_2}(a_3)$. By directedness, we can pick

$k \in I$ with $j_1 \leq k$ and $j_2 \leq k$. Now

$$\begin{aligned} f_{i_1 k}(a_1) &= f_{j_1 k}(f_{i_1 j_1}(a_1)) \\ &= f_{j_1 k}(f_{i_2 j_1}(a_2)) \\ &= f_{i_2 k}(a_2) \\ &= f_{j_2 k}(f_{i_2 j_2}(a_2)) \\ &= f_{j_2 k}(f_{i_3 j_2}(a_3)) \\ &= f_{i_3 k}(a_3), \end{aligned}$$

so $a_1 \sim a_3$.

The domain of the direct limit $\varinjlim M_i$ is $M_*/{\sim}$, whose elements have the form $[a]$ with $a \in M_*$. For each $i \in I$, there is a function $g_i \colon M_i \to \varinjlim M_i$ given by $g_i(a) = [a]$.

Given a tuple $b = (b_1, \ldots, b_n)$ from $\varinjlim M_i$, a **representative** of b is a tuple $a = (a_1, \ldots, a_n)$ from some fixed M_i such that $[a_j] = b_j$ for all $1 \leq j \leq n$, i.e., such that $g_i(a) = b$. If a is a representative of b, we write $b = [a]$. Remember that this notation implies that a lives in some fixed M_i.

Lemma 4.3. *Every tuple $b = (b_1, \ldots, b_n)$ from $\varinjlim M_i$ has a representative. Moreover, if $a = (a_1, \ldots, a_n)$ from M_i and $a' = (a'_1, \ldots, a'_n)$ from $M_{i'}$ are both representatives of b, then there exists $k \in I$ with $i \leq k$ and $i' \leq k$ such that $f_{ik}(a) = f_{i'k}(a')$.*

Proof. For each $1 \leq j \leq n$, $b_j = [c_j]$ for some $c_j \in M_{i_j}$. By directedness, there exists $k \in I$ with $i_j \leq k$ for all $1 \leq j \leq n$ (when $n = 0$, we use the fact that I is non-empty). Let $a_j = f_{i_j k}(c_j)$. Then $f_{i_j i_j}(a_j) = f_{i_j k}(c_j)$, so $a_j \sim c_j$, and $b_j = [c_j] = [a_j]$ for all $1 \leq j \leq n$. Thus $a = (a_1, \ldots, a_n)$ is a representative of b.

Now suppose $a = (a_1, \ldots, a_n)$ from M_i and $a' = (a'_1, \ldots, a'_n)$ from $M_{i'}$ are both representatives of b. For each $1 \leq j \leq n$, $[a_j] = b = [a'_j]$, so $a_j \sim a'_j$. Thus there exists k_j with $i \leq k_j$ and $i' \leq k_j$ such that $f_{ik_j}(a_j) = f_{i'k_j}(a'_j)$. By directedness, there exists $k \in I$ with $k_j \leq k$ for all $1 \leq j \leq n$. Now for each j, $f_{ik}(a_j) = f_{k_j k}(f_{ik_j}(a_j)) = f_{k_j k}(f_{i'k_j}(a'_j)) = f_{i'k}(a'_j)$. So $f_{ik}(a) = f_{i'k}(a')$, as desired. □

Lecture 8: 10/1

Similarly, for a finite variable context x, if $b \in (\varinjlim M_i)^x$, a **representative** of b is $a \in M_i^x$ for some i such that $g_i(a) = b$, i.e., $b = g_i \circ a$, and in this case we write $b = [a]$. Lemma 4.3 applies just as well to finite variable assignments as it does to finite tuples. We can rephrase Lemma 4.3 by writing that for finite x, $(\varinjlim M_i)^x$ is in natural bijection with $\varinjlim (M_i^x)$.

We make $\varinjlim M_i$ into an $\mathcal{L}$-structure as follows:

- For each n-ary function symbol $f \in \mathcal{L}$ and any n-tuple b in $\varinjlim M_i$ such that $b = [a]$ with a in M_i, $f^{\varinjlim M_i}(b) = [f^{M_i}(a)]$.
- For each n-ary relation symbol $R \in \mathcal{L}$ and any n-tuple b in $\varinjlim M_i$, $b \in R^{\varinjlim M}$ if and only if b has a representative a in some M_i such that $a \in R^{M_i}$.

We only need to check that the interpretations of the function symbols are well-defined. By Lemma 4.3, b has some representative a, so the definition makes sense. For well-definedness, suppose a in M_i and a' in $M_{i'}$ are both representatives for b. Then by Lemma 4.3, there is some $k \in I$ with $i \leq k$ and $i' \leq k$ such that $f_{ik}(a) = f_{i'k}(a)$. Since these maps are homomorphisms, we have $f_{ik}(f^{M_i}(a)) = f^{M_k}(f_{ik}(a)) = f^{M_k}(f_{i'k}(a')) = f_{i'k}(f^{M_{i'}}(a'))$. Thus $[f^{M_{i'}}(a')] = [f^{M_i}(a)]$, and $f^{\varinjlim M_i}$ is well-defined.

Lemma 4.4. *For all $i \in I$, the map $g_i \colon M_i \to \varinjlim M_i$ by $g_i(a) = [a]$ is a homomorphism, and for all $i \leq j$ in I, $g_j \circ f_{ij} = g_i$.*

Proof. Directly from the definition. For each n-ary function symbol f, since a is a representative of $g_i(a)$, $g(f^{M_i}(a)) = [f^{M_i}(a)] = f^{\varinjlim M_i}([a]) = f^{\varinjlim M_i}(g_i(a))$. And for each n-ary relation symbol R, if $a \in R^{M_i}$, then since a is a representative for $g_i(a)$, $g_i(a) \in R^{\varinjlim M_i}$.

Given $i \leq j$ in I and $a \in M_i$, we have $f_{jj}(f_{ij}(a)) = f_{ij}(a)$, so $[a] = [f_{ij}(a)]$, i.e., $g_i(a) = g_j(f_{ij}(a))$. □

Recall that the positive existential fragment $\exists^+$ is the smallest fragment containing all atomic formulas and closed under $\wedge$, $\vee$, and $\exists$. We proved (Theorem 2.13) that homomorphisms preserve all $\exists^+$-formulas. We will now see that $\exists^+$ formulas are reflected in direct limits, in a certain sense.

Lemma 4.5. *Let $(M_i)_{i \in I}$ be a directed family with direct limit $M = \varinjlim M_i$. Let $\varphi(x)$ be an $\exists^+$-formula, and let $b \in M^x$. Then $M \models \varphi(b)$ if and only if there exists $i \in I$ and a representative $a \in M_i^x$ of b such that $M_i \models \varphi(a)$.*

Proof. One direction is clear: If $M_i \models \varphi(a)$, then since g_i is a homomorphism and $g_i(a) = b$, $M \models \varphi(b)$ by Theorem 2.13.

We prove the converse by induction on φ.

Suppose φ is atomic of the form $t_1 = t_2$. Pick any representative $a \in M_i^x$ of b. Since $M \models \varphi(b)$, $t_1^M(g_i(a)) = t_2^M(g_i(a))$. Since g_i is a homomorphism, $[t_1^{M_i}(a)] = [t_2^{M_i}(a)]$. Thus there exists $i \leq j \in I$ such that $f_{ij}(t_1^{M_i}(a)) = f_{ij}(t_2^{M_i}(a))$. Since f_{ij} is a homomorphism, $t_1^{M_j}(f_{ij}(a)) = t_2^{M_j}(f_{ij}(a))$. Thus $M_j \models \varphi(f_{ij}(a))$. Since $b = g_i(a) = g_j(f_{ij}(a))$, $f_{ij}(a)$ is a representative of b.

Now suppose φ is atomic of the form $R(t_1, \ldots, t_n)$. Pick any representative $a \in M_i^x$ of b. Let c be the tuple $(t_1^M(b), \ldots, t_n^M(b))$. Since g_i is a homomorphism, the tuple $d = (t_1^{M_i}(a), \ldots, t_n^{M_i}(a))$ is a representative of c.

Since $c \in R^M$, there is also a representative d' of c in some M_j such that $d' \in R^{M_j}$. By Lemma 4.3, there exists $k \in I$ with $i \leq k$ and $j \leq k$ such that $f_{ik}(d) = f_{jk}(d')$. Call this tuple e. Since f_{jk} is a homomorphism, $e = f_{jk}(d') \in R^{M_k}$. And since f_{ik} is a homomorphism, $e = f_{ik}(d) = (t_1^{M_k}(f_{ik}(a)), \ldots, t_n^{M_k}(f_{ik}(a)))$. It follows that $M_k \models \varphi(f_{ik}(a))$, and since $f_{ik}(a)$ is a representative of b in M_k, we are done.

When φ is $\top$ or $\bot$, we can pick any $i \in I$ and any representative $a \in M_i^x$.

Suppose φ is $\psi \wedge \chi$. If $M \models \varphi(b)$, then $M \models \psi(b)$ and $M \models \chi(b)$. By induction we can pick representatives $a \in M_i^x$ and $a' \in M_{i'}^x$ such that $M_i \models \psi(a)$

and $M_{i'} \models \chi(a')$. By Lemma 4.3, we can pick $j \in I$ with $i \leq j$ and $i' \leq j$ such that $f_{ij}(a) = f_{i'j}(a')$. Call this assignment c. Since f_{ij} and $f_{i'j}$ are homomorphisms, $M_j \models \psi(c) \wedge \chi(c)$, and since c is a representative of b in M_j, we are done.

Suppose φ is $\psi \vee \chi$. If $M \models \varphi(b)$, then $M \models \psi(b)$ or $M \models \chi(b)$. In the first case, by induction we can pick a representative $a \in M_i^x$ such that $M_i \models \psi(a)$. Then $M_i \models \varphi(a)$, and we are done. The second case is similar.

Finally, suppose φ is $\exists y\, \psi$. If $M \models \varphi(b)$, then there exists $b' \in M$ such that $M \models \psi(b, b')$. By induction, we can pick a representative $(a, a') \in M_i^{xy}$ such that $M_i \models \psi(a, a')$. Thus $M_i \models \varphi(a)$, and we are done. □

As usual for our fragments, an $\exists^+$-theory is a set of sequents of the form $\varphi \vdash_x \psi$, where φ and ψ are $\exists^+$-formulas in context x.

Theorem 4.6. *Let T be an $\exists^+$-theory. Given a directed family $(M_i)_{i\in I}$ of models of T, we have $\varinjlim M_i \models T$.*

Proof. Assume each M_i is a model of T, and let $M = \varinjlim M_i$. Let $\varphi \vdash_x \psi$ be a sequent in T. To show $M \models \varphi \vdash_x \psi$, let $b \in M^x$, and assume $M \models \varphi(b)$. By Lemma 4.5, there exists $i \in I$ and $a \in M_i^x$ with $g_i(a) = b$ such that $M_i \models \varphi(a)$. Since $M_i \models T$, $M_i \models \psi(a)$. Now since g_i is a homomorphism, by Theorem 2.13, $M \models \psi(b)$. □

This theorems shows that all the classes of structures considered in Section 3.2 are closed under direct limits (groups, torsion-free groups, R-modules, rings, reduced rings, integral domains, and fields, as well as posets and linear orders, both in their strict and non-strict formulations).

4.2 Positively existentially closed structures

We have seen that homomorphisms preserve $\exists^+$-formulas. So if $f\colon M \to N$ is a homomorphism and $a \in M^x$, $f(a)$ satisfies all the $\exists^+$-formulas that a does, and possibly more. We now define the structures M which are maximal from this perspective, in the sense that any $\exists^+$-formula that could be satisfied by $f(a)$ for some homomorphism f is already satisfied by a.

Definition 4.7. Let $\mathcal{K}$ be a class of $\mathcal{L}$-structures (usually the class of models of a theory T). We say that a structure $M \in \mathcal{K}$ is **positively existentially closed** (or **$\exists^+$-closed**, or **PEC**) in $\mathcal{K}$, if for every homomorphism $f\colon M \to N$ with $N \in \mathcal{K}$, f reflects $\exists^+$-formulas.

Example 4.8. In the language $\mathcal{L}_{\mathrm{SOrd}} = \{<\}$, a PEC linear order $(L, <)$ must be:

- Nonempty. We can always find a non-empty linear order L' with $L \subseteq L'$. The inclusion map reflects the $\exists^+$-sentence $\exists x\, \top$.

- Unbounded above. For any $a \in L$, we can always find a linear order L' with $L \subseteq L'$ such that L' contains a new element greater than a. The inclusion map reflects the $\exists^+$-formula $\exists y\,(y > x)$, so a is not a greatest element in L.
- Unbounded below, similarly.
- Dense. For any $a < b$ in L, we can always find a linear order L' with $L \subseteq L'$ such that L' contains a new element strictly between a and b. The inclusion map reflects the $\exists^+$ formula $\exists z\,(x < z \wedge z < y)$, so L contains an element strictly between a and b.

In fact, a linear order is PEC if and only if it is dense with no greatest or least element (e.g. $(\mathbb{Q}, <)$ and $(\mathbb{R}, <)$ are PEC linear orders). We will justify this claim later.

In the language $\mathcal{L}_{\text{Ord}} = \{\leq\}$, a PEC linear order $(L, \leq)$ has only one element. Indeed, L admits a homomorphism $f\colon L \to \mathbf{1}$, which reflects the formula $x = y$. Since $f(a) = f(b)$ for all $a, b \in L$, we have $a = b$.

Example 4.9. A PEC ring is trivial. Indeed, since every ring R admits a homomorphism f to the zero ring Z, and $Z \models 0 = 1$, a PEC ring must satisfy $0 = 1$.

However, I claim that a commutative ring which is PEC in the class of non-zero rings is an algebraically closed field. Suppose R is PEC in the class of non-zero rings. Let $M \subseteq R$ be a maximal ideal, and let $q\colon R \to R/M$ be the quotient homomorphism. Let $\varphi(x)$ be the $\exists^+$ formula $x = 0 \vee \exists y\,(x \cdot y = 1)$. Then for any $a \in R$, $R/M \models \varphi(a)$, so $R \models \varphi(a)$. Thur R is a field. Now for any non-constant polynomial $p \in R[x]$, we can embed R in a field K containing a root of p (e.g., $K = R[x]/(q)$, where q is an irreducible factor of p). Since $K \models \exists x\, p(x) = 0$, $R \models \exists x\, p(x) = 0$, so R is algebraically closed.

In fact, every algebraically closed field is PEC in the class of non-trivial rings. We will justify this claim later.

Lecture 9: 10/6

Theorem 4.10. *Let T be an $\exists^+$-theory. For any $M \models T$, there exists a homomorphism $f\colon M \to N$ where N is PEC in the class of models of T.*

First, we prove that any model M admits a homomorphism to a model N which satisfies the PEC property, but only for assignments in the image of M.

Lemma 4.11. *Let T be an $\exists^+$-theory. For any $M \models T$, there exists a homomorphism $f\colon M \to N$, where $N \models T$ and has the following property: for any homomorphism $g\colon N \to N'$ with $N' \models T$, any $\exists^+$-formula $\varphi(x)$, and any $a \in M^x$, if $N' \models \varphi(g(f(a)))$, then $N \models \varphi(f(a))$.*

Proof. For some cardinal κ, enumerate as $(\varphi_\alpha(a_\alpha))_{\alpha<\kappa}$ the set of all $\varphi(a)$, where $\varphi(x)$ is an $\exists^+$-formula and $a \in M^x$.

We build a κ-indexed directed family $(M_\alpha)_{\alpha\leq\kappa}$ of models of T by transfinite recursion on κ.

Base case: Let $M_0 = M$ and $f_{00} = \mathrm{id}_{M_0}$.

Inductive step: Given M_α, if there exists some $N \models T$ and a homomorphism $h\colon M_\alpha \to N$ such that $N \models \varphi_\alpha(h(f_{0\alpha}(a_\alpha)))$, let $M_{\alpha+1} = N$ and $f_{\alpha(\alpha+1)} = h$. Otherwise, if there is no such pair (N, h), set $M_{\alpha+1} = M_\alpha$ and $f_{\alpha(\alpha+1)} = \mathrm{id}_{M_\alpha}$. Now set $f_{(\alpha+1)(\alpha+1)} = \mathrm{id}_{M_{\alpha+1}}$, and for all $\beta < \alpha$, set $f_{\beta(\alpha+1)} = f_{\alpha(\alpha+1)} \circ f_{\beta\alpha}$. Then $(M_\beta)_{\beta \leq \alpha+1}$ is a directed family of models of T.

Limit step: Suppose $\gamma \leq \kappa$ is a limit ordinal. By recursion we have constructed $(M_\alpha)_{\alpha<\gamma}$, a directed family of models of T. Let $M_\gamma = \varinjlim M_\alpha$. By Theorem 4.6, $M_\gamma \models T$. Let $f_{\gamma\gamma} = \mathrm{id}_{M_\gamma}$, and for all $\alpha < \gamma$, let $f_{\alpha\gamma} = g_\alpha$, the canonical homomorphism to the direct limit. Then $(M_\alpha)_{\alpha \leq \gamma}$ is again a directed family of models of T.

Finally, let $N = M_\kappa$ and $f = f_{0\kappa}$. Suppose $g\colon N \to N'$ is a homomorphism with $N' \models T'$ and $N' \models \varphi(g(f(a)))$ for some $\exists^+$-formula $\varphi(x)$ and some $a \in M^x$. Then $\varphi(a)$ is $\varphi_\alpha(a_\alpha)$ for some $\alpha < \kappa$.

Note that $g \circ f = g \circ f_{0\kappa} = (g \circ f_{\alpha\kappa}) \circ f_{0\alpha}$, so $h = g \circ f_{\alpha\kappa}\colon M_\alpha \to N'$ is a homomorphism such that $N' \models \varphi_\alpha(h(f_{0\alpha}(a_\alpha)))$. Because of the existence of the pair (N', h), we must have $M_{\alpha+1} \models \varphi_\alpha(f_{\alpha(\alpha+1)}(f_{0\alpha}(a_\alpha)))$. Since the homomorphism $f_{(\alpha+1)\kappa}$ preserves $\exists^+$-formulas, and $f_{(\alpha+1)\kappa} \circ f_{\alpha(\alpha+1)} \circ f_{0\alpha} = f_{0\kappa} = f$, we have $N \models \varphi(f(a))$, as desired. □

Proof of Theorem 4.10. Suppose $M \models T$, an $\exists^+$-theory. We build an ω-indexed directed family $(M_n)_{n<\omega}$ of models of T by recursion.

Base case: Let $M_0 = M$ and $f_{00} = \mathrm{id}_{M_0}$.

Inductive step: Given M_n, let $M_{n+1} \models T$ and $f_{n(n+1)}\colon M_n \to M_{n+1}$ be the model and homomorphism provided by Lemma 4.11. Now let $f_{(n+1)(n+1)} = \mathrm{id}_{M_{n+1}}$ and for all $m < n$, let $f_{m(n+1)} = f_{n(n+1)} \circ f_{mn}$.

Having constructed the directed family $(M_n)_{n<\omega}$, let $N = \varinjlim (M_n)_{n<\omega}$, and let $f = g_0$, the canonical map into the direct limit. By Theorem 4.6, $N \models T$. It remains to show that N is PEC.

Suppose $h\colon N \to N'$ is a homomorphism and $N' \models T$. Let $\varphi(x)$ be an $\exists^+$-formula, let $a \in N^x$, and assume $N' \models \varphi(h(a))$. By 4.3, a has a representative $a' \in M_n^x$ for some $n < \omega$. We can factor $g_n = g_{n+1} \circ f_{n(n+1)}$, so $h(a) = h(g_n(a')) = h(g_{n+1}(f_{n(n+1)}(a')))$. By the partial PEC property of M_{n+1}, since $N' \models \varphi(h(g_{n+1})(f_{n(n+1)}(a')))$ and $h \circ g_{n+1}\colon M_{n+1} \to N'$ is a homomorphism to a model of T, we have $M_{n+1} \models \varphi(f_{n(n+1)}(a'))$. And since the homomorphism g_{n+1} preserves the $\exists^+$-formula φ, we have $N \models \varphi(a)$, as desired. □

4.3 Compactness for $\exists^+$-theories

We begin with a "normal form" lemma for the $\exists^+$-fragment. A **positive primitive** (p.p.) formula is one of the form $\exists y \bigwedge_{i=1}^n \varphi_i(x, y)$, where each φ_i is atomic. Here y is a finite set of variables, possibly empty (so a conjunction of atomic formulas is p.p.).

Lemma 4.12. *Every $\exists^+$-formula is logically equivalent to a disjunction of p.p. formulas.*

Proof. Let φ be an $\exists^+$-formula. We proceed by induction.

If φ is atomic, it is a singleton disjunction of the empty quantification of a singleton conjunction of atomic formulas.

If φ is $\bot$, it is the empty disjunction of p.p. formulas.

If φ is $\top$, it is the singleton disjunction of the empty quantification of the empty conjunction of atomic formulas.

Suppose φ is $\psi \vee \chi$. By induction, we may assume ψ and χ are disjunctions of p.p. formulas. Then φ is too.

Suppose φ is $\psi \wedge \chi$. By induction, we may assume ψ and χ are disjunctions of p.p. formulas. Distributing the $\wedge$ over the disjunctions, φ is equivalent to a disjunction of formulas, each of which is the conjunction of two p.p. formulas. Up to renaming bound variables, we may assume the quantified variables in each of the p.p. formulas are disjoint. Now when y and z are disjoint, $(\exists y\, \theta_1) \wedge (\exists z\, \theta_2)$ is equivalent to $\exists y \exists z\, (\theta_1 \wedge \theta_2)$, which is p.p. when θ_1 and θ_2 are conjunctions of atomic formulas.

Suppose φ is $\exists y\, \psi$. By induction, we may assume ψ is a disjunction of p.p. formulas, $\bigvee_{i=1}^n \psi_i$. Now $\exists y \bigvee_{i=1}^n \psi_i$ is equivalent to $\bigvee_{i=1}^n \exists y\, \psi_i$, and each formula $\exists y\, \psi_i$ is p.p. $\square$

Lecture 10: 10/8

Let T be an $\exists^+$-theory. We say that a sequent S is **finitely entailed** by T if there is some finite $T_0 \subseteq T$ such that $T_0 \models S$.

Lemma 4.13. *Let T be an $\exists^+$-theory. Let T_C be the set of all constraint clauses that are finitely entailed by T. If M is PEC in the class of models of T_C, then $M \models T$.*

Proof. To show $M \models T$, it suffices to show that $M \models (\varphi \vdash_x \psi)$ for every sequent $(\varphi \vdash_x \psi) \in T$.

Let $a \in M^x$, and assume for contradiction that $M \models \varphi(a)$ and $M \not\models \psi(a)$. By Lemma 4.12, φ is equivalent in M to a disjunction of p.p. formulas. Since a satisfies φ in M, a satisfies one of the disjuncts $\exists y\, \varphi'$ in M. Then there exists $b \in M^y$ such that $M \models \varphi'(a, b)$. Note that $\varphi'(x, y)$ is a conjunction of atomic formulas, and $\varphi' \vdash_{xy} \varphi$ is a validity (true in every structure).

Similarly, by Lemma 4.12, ψ is equivalent to $\bigvee_{i=1}^n \exists z^i\, \psi_i(x, z^i)$, where each ψ_i is a conjunction of atomic formulas. Then for all $1 \leq i \leq n$, $M \not\models \exists z^i\, \psi_i(a, z^i)$.

Fixing an i, let c^i be a set of new constant symbols, one for each variable in z^i, and consider the $\mathcal{L}(M \cup c^i)$-theory

$$\Sigma_i = T_C \cup \mathrm{Diag}^+(M) \cup \{\psi_i(a, c^i)\}.$$

Note that Σ_i is a Horn theory with constraints (in fact, it is equivalent to a set of atomic axioms, together with a set of constraint clauses). If Σ_i has a model N, then because $N \models \mathrm{Diag}^+(M)$, there is a homomorphism $h\colon M \to N|_{\mathcal{L}}$ with $h(m) = m^N$, and since $N \models \psi_i(a, c^i)$, $N \models \exists z^i\, \psi_i(h(a), z^i)$ (with $(c^i)^N$ witnessing the existential quantifiers). Since $N \models T_C$ and M is a PEC model of T_C, $M \models \exists z^i\, \psi_i(a, z^i)$, contradiction.

Thus Σ_i has no model. By Theorem 3.14, there is a finite set $\Delta_i \subseteq \mathrm{Diag}^+(M)$ and a constraint clause $C_i \in T_C$ such that $\{C_i\} \cup \Delta_i \cup \{\psi_i(a, c^i)\}$ has no model.

Since C is finitely entailed by T, there is some finite $T_i \subseteq T$ such that $T_i \models C$, so $T_i \cup \Delta_i \cup \{\psi_i(a, c^i)\}$ has no model.

Having found Δ_i and T_i for each $1 \leq i \leq n$, let $\Delta = \bigcup_{i=1}^n \Delta_i$, and let $T^* = \{\varphi \vdash_x \psi\} \cup \bigcup_{i=1}^n T_i \subseteq_{\text{fin}} T$. Note that $\Delta \subseteq_{\text{fin}} \text{Diag}^+(M)$. Write the conjunction of the atomic sentences in Δ as $\delta(a, b, m)$, separating the constants naming the elements of the assignments a and b from the others in m.

Let C^* be the constraint clause $\varphi'(x, y) \wedge \delta(x, y, w) \vdash_{xyw} \bot$. I claim that $T^* \models C^*$, so that C^* is finitely entailed by T. Let $M' \models T^*$, let $a', b', m' \in (M')^{xyw}$, and suppose $M' \models \varphi'(a', b') \wedge \delta(a', b', m')$. Since $\varphi' \vdash_{xy} \varphi$ is a validity and $M' \models (\varphi \vdash_x \psi)$, we have $M' \models \psi(a')$. Then there is some $1 \leq i \leq n$ such that $M' \models \exists z^i\, \psi_i(a', z^i)$, so there is $c' \in (M')^{z^i}$ such that $M' \models \psi_i(a', c')$.

Since $M' \models T^*$, $M' \models T_i$. Interpreting the constant symbols a, b, c^i, m as a', b', c', m' in M', $M' \models \psi_i(a, c^i)$, and since $M' \models \delta(a', b', m')$, $M' \models \Delta_i$. This contradicts the fact that $T_i \cup \Delta_i \cup \{\psi_i(a, c^i)\}$ has no model.

Thus $T^* \models C^*$, and hence $C^* \in T_C$. But $M \models \varphi'(a, b) \wedge \delta(a, b, m)$, so $M \not\models C^*$, contradicting $M \models T_C$. Thus $M \models (\varphi \vdash_x \psi)$, as desired. □

Theorem 4.14. *Let T be an $\exists^+$-theory. Assume that every finite subtheory $T_0 \subseteq T$ has a model. Then T has a model.*

Proof. Let T_C be the set of all constraint clauses that are finitely entailed by T. I claim that T_C has a model. By compactness for Horn theories with constraints (Theorem 3.14), it suffices to show that every single constraint clause $C \in T_C$ has a model. Since C is finitely entailed by T, there is some finite $T_0 \subseteq T$ such that $T_0 \models C$. By our hypothesis, T_0 has a model, and thus C has a model M. [In fact, the initial model of the Horn part of T_C is a model of T_C. Since the Horn part is empty, the closed term algebra is a model of T_C.]

Now T_C is an $\exists^+$-theory, so by Theorem 4.10, T_C has a PEC model M. By Lemma 4.13, $M \models T$. □

Corollary 4.15. *Let T be an $\exists^+$-theory, and let $S = \varphi \vdash_x \psi$ be an $\exists^+$-sequent. Then S is finitely entailed by T if and only if S is entailed by T.*

Proof. If there exists $T_0 \subseteq_{\text{fin}} T$ such that $T_0 \models S$, then clearly $T \models S$.

Conversely suppose $T \models S$. Let c be a set of new constant symbols, one for every variable in x, and consider the $\mathcal{L}(c)$-theory T':

$$T \cup \{\top \vdash_\varnothing \varphi(c), \psi(c) \vdash_\varnothing \bot\}.$$

Then T' has no model, since if $M \models T'$, then $M \models T$, $M \models \varphi(c^M)$, and $M \not\models \psi(c^M)$, contradicting $T \models S$.

By Theorem 4.14, there is a finite $T_0 \subseteq T$ such that

$$T_0 \cup \{\top \vdash_\varnothing \varphi(c), \psi(c) \vdash_\varnothing \bot\}$$

has no model. It follows that $T_0 \models \varphi \vdash_x \psi$. Indeed, if $M \models T_0$ and $a \in M^x$ such that $M \models \varphi(a)$ but $M \models \neg\psi(a)$, then we can expand M to an $\mathcal{L}(c)$-structure by $c^M = a$, and M is a model of $T_0 \cup \{\top \vdash_\varnothing \varphi(c), \psi(c) \vdash_\varnothing \bot\}$, contradiction. □

Corollary 4.16. *Let T be an $\exists^+$-theory, and let T_C be the set of all constraint clauses entailed by T. If M is a PEC model of T_C, then $M \models T$.*

Proof. By Corollary 4.15, T_C is equal to the set of constraint clauses *finitely* entailed by T. Then the result follows from Lemma 4.13. □

4.4 Morleyization

We now bootstrap our compactness theorem for $\exists^+$-theories up to arbitrary first-order theories. This may seem like a heavy lift, but it turns out that the full first-order fragment can be interpreted, in an appropriate sense, in the $\exists^+$-fragment. We give the construction in greater generality.

We say a fragment $\mathcal{F}$ is **$\forall$-free** if no formula in $\mathcal{F}$ contains the universal quantifier $\forall$. For example, the literal fragment Lit and the existential fragment $\exists$ are $\forall$-free. The elementary fragment FO is not $\forall$-free, but we can let FO* be the set of all first-order formulas with no universal quantifiers. Every formula in FO is logically equivalent to one in FO*, by rewriting $\forall$ as $\neg\exists\neg$, so the $\forall$-free fragment FO* is just as expressive as FO.

For the rest of the construction, we fix a $\forall$-free fragment $\mathcal{F}$. Recall that we assume fragments are closed under subformula and substitution of terms for free variables.

Fix also a countable set of variables $V = \{x_n \mid n \in \mathbb{N}\}$. Since every formula in $\mathcal{F}$ mentions only finitely many variables, at the expense of renaming variables, we can focus on formulas with variable contexts from V.

For each formula $\varphi(x) \in \mathcal{F}$ where $x = \{x_0, \ldots, x_{n-1}\}$ for some n, introduce a new n-ary relation symbol R_φ. Let $\mathcal{L}_\mathcal{F} = \mathcal{L} \cup \{R_\varphi \mid \varphi(x) \in \mathcal{F}\}$.

For each such formula $\varphi(x) \in \mathcal{F}$, we now introduce some new $\exists^+$-axioms in the language $\mathcal{L}_\mathcal{F}$, depending on the form of $\varphi(x)$. In the recursion, we use the fact that $\mathcal{F}$ is closed under subformula.

- If φ is atomic: $R_\varphi(x) \vdash_x \varphi$ and $\varphi \vdash_x R_\varphi(x)$.
- If φ is $\top$: $\top \vdash_x R_\varphi(x)$.
- If φ is $\bot$: $R_\varphi(x) \vdash_x \bot$.
- If φ is $\psi \wedge \chi$: $R_\varphi(x) \vdash_x R_\psi(x) \wedge R_\chi(x)$ and $R_\psi(x) \wedge R_\chi(x) \vdash_x R_\varphi(x)$.
- If φ is $\psi \vee \chi$: $R_\varphi(x) \vdash_x R_\psi(x) \vee R_\chi(x)$ and $R_\psi(x) \vee R_\chi(x) \vdash_x R_\varphi(x)$.
- If φ is $\neg\psi$: $R_\varphi(x) \wedge R_\psi(x) \vdash_x \bot$ and $\top \vdash_x R_\varphi(x) \vee R_\psi(x)$.
- If φ is $\exists y\, \psi$: $R_\varphi(x) \vdash_x \exists y\, R_\psi(x, y)$ and $\exists y\, R_\psi(x, y) \vdash_x R_\varphi(x)$.

Let $T_\mathcal{F}$ be the set of axioms listed above, for every formula in $\mathcal{F}$ with variable context $x = \{x_0, \ldots, x_n\}$ for some n. Note that $T_\mathcal{F}$ is an $\exists^+$-theory.

Given an $\mathcal{F}$-theory T in $\mathcal{L}$, let $\widehat{T}$ be the $\exists^+$ theory:

$$T_\mathcal{F} \cup \{R_\varphi(x) \vdash_x R_\psi(x) \mid \varphi \vdash_x \psi \in T\}.$$

Given an $\mathcal{L}$-structure M, it is easy to prove by induction that for any M has a unique expansion to a model $\widehat{M} \models T_{\mathcal{F}}$, in which $R_{\varphi}^{\widehat{M}} = \{a \in M^x \mid M \models \varphi(a)\}$. Moreover, $M \models T$ if and only if $\widehat{M} \models \widehat{T}$.

Lecture 11: 10/13

Let $\mathcal{F}$ be a $\forall$-free fragment. We have seen above that there is a bijection between $\mathcal{L}$-structures M and $\mathcal{L}_{\mathcal{F}}$-structures $\widehat{M}$ which are models of $T_{\mathcal{F}}$. Additionally, for a function $h\colon M \to N$, h is an $\mathcal{F}$-morphism if and only if $h\colon \widehat{M} \to \widehat{N}$ is a homomorphism. That is, we have an isomorphism of categories between the category of $\mathcal{L}$-structures and $\mathcal{F}$-morphisms and the category of models of $T_{\mathcal{F}}$ and homomorphisms. This allows us to lift our results about direct limits along directed families with homomorphisms to directed families with stronger connecting maps.

- Let $(M_i)_{i\in I}$ be a directed family such that every map $f_{ij}\colon M_i \to M_j$ is an $\mathcal{F}$-morphism. Let $M = \varinjlim M_i$. There is a corresponding directed family $(\widehat{M_i})_{i\in I}$. Since each $\widehat{M_i} \models T_{\mathcal{F}}$, by Theorem 4.6, $\varinjlim M_i \models T_{\mathcal{F}}$. Since M has a unique expansion to a model of $T_{\mathcal{F}}$, $\varinjlim \widehat{M_i} = \widehat{M}$.

- The canonical map $g_i\colon \widehat{M_i} \to \widehat{M}$ is a homomorphism by Lemma 4.4, so $g_i\colon M_i \to M$ is an $\mathcal{F}$-morphism.

- Let $\varphi(x)$ be an $\exists^+(\mathcal{F})$-formula, and let $b \in M^x$. Then $\varphi(x)$ translates to an $\exists^+$-formula $\widehat{\varphi}(x)$ in $\mathcal{L}_{\mathcal{F}}$. We have $M \models \varphi(b)$ if and only if $\widehat{M} \models \widehat{\varphi}(b)$ if and only if there exists $i \in I$ and a representative $a \in \widehat{M_i}^x$ of b such that $\widehat{M_i} \models \widehat{\varphi}(a)$ (by Lemma 4.5) if and only if $M_i \models \varphi(a)$,.

- Let T be an $\exists^+(\mathcal{F})$-theory. We can translate T into an $\exists^+$-theory $\widehat{T}$ in $\mathcal{L}_{\mathcal{F}}$. If $M_i \models T$ for all $i \in I$, then $\widehat{M_i} \models \widehat{T}$ for all $i \in I$, so $\widehat{M} \models \widehat{T}$ (by Theorem 4.6), and $M \models T$.

For example, take $\mathcal{F}$ to be the literal fragment Lit, so $\mathcal{F}$-morphisms are embeddings and $\exists^+(\mathcal{F})$ is the existential fragment $\exists$. Then given a directed family with all connecting maps embeddings, each of the structures in the family embeds in the direct limit, and the direct limit preserves satisfaction of $\exists$-theories.

For another example, take $\mathcal{F}$ to be the $\forall$-free fragment FO^*, and note that $\mathcal{F}$-morphisms are elementary embeddings and $\exists^+(\mathcal{F}) = \mathrm{FO}^*$. Given a directed family with all connecting maps elementary embeddings, each of the structures in the family embeds elementarily in the direct limit, and the direct limit preserves satisfaction of arbitrary first-order theories.

Example 4.17. We end with an example demonstrating the necessity of the "$\forall$-free" condition in the construction above.

Consider the empty language $\mathcal{L} = \varnothing$, and the fragment $\mathcal{F}$ obtained by closing the atomic formulas under universal quantifiers.

For each $i \in \omega$, let $A_i = \{a_{i0}, a_{i1}\}$, where the a_{ij} are pairwise distinct. For all $i < j$, let $f_{ij}\colon A_i \to A_j$ be the constant map sending both elements of A_i to a_{j0}. It is easy to check that the f_{ij} are $\mathcal{F}$-morphisms.

The direct limit $A = \varinjlim A_i$ has a single point, since for any a_{ij} and $a_{i'j'}$ with $i \leq i'$, we have $f_{i(i'+1)}(a_{ij}) = f_{i'(i'+1)}(a_{i'j'}) = a_{(i'+1)0}$. Then A satisfies the $\mathcal{F}$-sentence $\varphi\colon \forall x \forall y\, (x = y)$, but there is no $i \in \omega$ such that $A_i \models \varphi$, so the direct limit does not reflect $\mathcal{F}$-formulas in the sense of Lemma 4.5.

Similarly, each A_i is a model of the $\mathcal{F}$-sentence $\forall x \forall y (x = y) \vdash_{\varnothing} \bot$, but A not does satisfy this sentence. Thus the direct limit does not preserve $\mathcal{F}$-theories in the sense of Theorem 4.6.

5 First-order compactness

5.1 The compactness theorem and some first consequences

Theorem 5.1 (Compactness theorem)**.** *Let T be a first-order theory. If every finite subset of T has a model, then T has a model.*

Proof. T has no model. We will show that a finite subset of T has no model.

Let FO^* be the $\forall$-free fragment of first-order logic. Then we can rewrite T to an equivalent FO^*-theory T^*. Since T has no model, the $\exists^+$-theory $\widehat{T^*}$ has no model. By compactness for $\exists^+$-theories (Theorem 4.14), there is a finite subset $T_0 \subseteq \widehat{T^*}$ such that T_0 has no model. Then $T_0 \subseteq T_0 \cup T_{\mathrm{FO}^*} = \widehat{T_1}$ for some $T_1 \subseteq_{\mathrm{fin}} T^*$, and $\widehat{T_1}$ has no model. Then T_1 has no model (otherwise, if $M \models T_1$, then $\widehat{M} \models \widehat{T_1}$). But T_1 is equivalent to a finite subset of T, so we are done. □

We will see three kinds of applications of the compactness theorem in this section: (1) building structures with specified properties, (2) showing the limitations of first-order logic: that some classes of structures are not axiomatizable and some sets and relations are not definable, and (3) transferring theorems from finite cases to the general case.

Definition 5.2. A theory T is **satisfiable** if T has a model. A theory T is **complete** if it is satisfiable, and for every sentence φ, either $T \models \varphi$ or $T \models \neg\varphi$.

Example 5.3. Let A be any structure. Then the **complete theory of** A is

$$\mathrm{Th}(A) = \{\varphi \text{ an } \mathcal{L}\text{-sentence} \mid A \models \varphi\}.$$

The name is justified, since for any sentence φ, either $A \models \varphi$ or $A \not\models \varphi$. In the first case, $\varphi \in \mathrm{Th}(A)$, while in the second case $A \models \neg\varphi$, so $\neg\varphi \in \mathrm{Th}(A)$.

Definition 5.4. Structure A and B are **elementarily equivalent**, written $A \equiv B$, if $\mathrm{Th}(A) = \mathrm{Th}(B)$, i.e., for all sentences φ, $A \models \varphi$ if and only if $B \models \varphi$.

Elementarily equivalent structures are indistinguishable from the perspective of sentences of first-order logic. Note that if $h\colon A \to B$ is an elementary embedding, then $A \equiv B$, since h preserves and reflects the truth of all sentences. In particular, isomorphic structures are elementarily equivalent.

Example 5.5. The language of arithmetic is $\mathcal{L}_{\mathrm{Arith}} = \{<, 0, 1, +, \times\}$. Consider the $\mathcal{L}_{\mathrm{Arith}}$-structure $\mathbb{N} = (\mathbb{N}; \leq, 0, 1, +, \times)$. The complete theory $\mathrm{Th}(\mathbb{N})$ is called **true arithmetic**, and its model $\mathbb{N}$ is the **standard model** of arithmetic. Let's use compactness to show that $\mathrm{Th}(\mathbb{N})$ also has *nonstandard* models.

Let $\mathcal{L} = \mathcal{L}_{\mathrm{Arith}} \cup \{c\}$, where c is a new constant symbol. For each $n \in \mathbb{N}$, let $\overline{n}$ be the term $\underbrace{1 + 1 + \cdots + 1}_{n \text{ times}}$. Define $T = \mathrm{Th}(\mathbb{N}) \cup \{c \neq \overline{n} \mid n \in \mathbb{N}\}$.

By compactness, to show that T is satisfiable, it suffices to show that any finite subset is satisfiable. For any finite subset $T_0 \subseteq_{\mathrm{fin}} T$, pick some $m \in \mathbb{N}$

such that $(c \neq \overline{m}) \notin T_0$. Then we can expand $\mathbb{N}$ to a model of T_0 by interpreting c as m.

Thus T has a model $(\mathcal{N}; \leq, 0, 1, +, \times, c)$. This model contains **standard** elements of the form $\overline{n}^{\mathcal{N}}$ for all $n \in \mathbb{N}$, but it also contains the element $c^{\mathcal{N}}$, which is not equal to any standard element. Since $\mathbb{N}$ is generated by $\varnothing$, and $\mathcal{N}$ satisfies all atomic sentences true in $\mathbb{N}$, there is a unique homomorphism $h\colon \mathbb{N} \to \mathcal{N}$, $n \mapsto \overline{n}^{\mathcal{N}}$. But h is not surjective, since $c^{\mathcal{N}}$ is not in its image. It follows that $\mathcal{N}$ is elementarily equivalent to $\mathbb{N}$ but not isomorphic to $\mathbb{N}$.

Example 5.6. Let $\mathcal{L}_{\mathbb{R},\text{full}}$ be the language containing an n-ary function symbol f for every function $f\colon \mathbb{R}^n \to \mathbb{R}$ (that includes a constant symbol r for every $r \in \mathbb{R}$) and an n-ary relation symbol R for every subset $R \subseteq \mathbb{R}^n$. We view $\mathbb{R}$ as an $\mathcal{L}_{\mathbb{R},\text{full}}$-structure in the obvious way.

Let $\mathcal{L} = \mathcal{L}_{\mathbb{R},\text{full}} \cup \{\varepsilon\}$, where ε is a new constant symbol. Define

$$T = \text{Th}(\mathbb{R}) \cup \{\varepsilon > 0\} \cup \{\varepsilon < r \mid r \in \mathbb{R}_{>0}\}.$$

By compactness, to show that T is satisfiable, it suffices to show that any finite subset is satisfiable. For any finite subset $T_0 \subseteq_{\text{fin}} T$, there is a least positive real number r such that $\varepsilon < r$ is in T_0. Then we can expand $\mathbb{R}$ to a model of T_0 by interpreting ε as $\frac{r}{2}$.

Thus T has a model $\mathcal{R}$. Since every element of $\mathbb{R}$ is named by a constant symbol in $\mathcal{L}_{\mathbb{R},\text{full}}$, $\text{Th}(\mathbb{R})$ is the elementary diagram of $\mathbb{R}$, so there is an elementary embedding $\mathbb{R} \to \mathcal{R}$, $r \mapsto r^{\mathcal{R}}$. But unlike $\mathbb{R}$, $\mathcal{R}$ contains an element $\varepsilon^{\mathcal{R}}$ which is infinitesimal: positive but smaller than every standard natural number.

It turns out that in models of T, naive calculus with infinitesimals becomes meaningful. For example, if a function $f\colon \mathbb{R} \to \mathbb{R}$ is differentiable at a, then $f'(a)$ is the unique real number that differs by an infinitesimal quantity from the evaluation of the term $\frac{f(a+\varepsilon)-f(a)}{\varepsilon}$ in $\mathcal{R}$ (note that we can make this definition without introducing the ε-δ definition of limit). This approach to calculus with infinitesimals is called **nonstandard analysis**.

Recall that the four color theorem in graph theory says that every finite planar graph admits a coloring by at most four colors. We show how to use compactness to automatically lift this theorem to the case of infinite graphs.

Example 5.7. Let G be an infinite planar graph. Let $\mathcal{L} = \mathcal{L}_{\text{Graph}}(G) \cup \{P_i \mid i < 4\}$, where the P_i are four new unary relation symbols. Define T to have the following axioms:

- $\text{Diag}_{\mathcal{L}_{\text{Graph}}}(G)$.
- $\forall x \bigvee_{i<4} P_i(x)$.
- $\forall x \bigwedge_{i<j<4} \neg(P_i(x) \wedge P_j(x))$.
- $\forall x \forall y\, (xEy \to \bigwedge_{i<4} \neg(P_i(x) \wedge P_i(y))$.

A finite subset $T_0 \subseteq T$ contains only a finite subset of $\text{Diag}(G)$. This subset describes a finite graph G_0. Note that G_0 is planar, since it is a subgraph of G. By the four color theorem, G_0 admits a coloring by at most four colors. We can expand G_0 to an $\mathcal{L}$-structure by interpreting the constant symbols naming elements of $G \setminus G_0$ arbitrarily, and interpreting the P_i as the colors. Then $G_0 \models T_0$.

By compactness, T has a model G', which is a graph with a four coloring. Note that $G'|_{\mathcal{L}_{\text{Graph}}}$ is not necessarily isomorphic to G. But since $G' \models \text{Diag}_{\mathcal{L}_{\text{Graph}}}(G)$, there is an embedding $h\colon G \to G'|_{\mathcal{L}_{\text{Graph}}}$, so G' has a subgraph isomorphic to G. Restricting the colors to this subgraphs gives a four coloring of G.

Lecture 12: 10/15

Finally, we show how compactness can be a tool for proving non-definability results.

Example 5.8. The language $\mathcal{L}_{\text{Graph}}$ consists of a binary relation symbol E. We will show that there is no first-order theory T such that $G \models T$ if and only if G is a connected graph.

Suppose for contradiction that such a theory T exists. Let $\mathcal{L} = \mathcal{L}_{\text{Graph}} \cup \{c, d\}$, where c and d are two new constant symbols, and let $T' = T \cup \{\varphi_n \mid n \in \omega\}$, where φ_0 is the sentence $(c \neq d)$, φ_1 is the sentence $\neg(cEd)$, and for $n \geq 2$, φ_n is the sentence expressing that there is no path of length n from c to d:

$$\neg\exists x_1 \ldots \exists x_{n-1} \left(cEx_1 \wedge x_{n-1}Ed \wedge \bigwedge_{i=1}^{n-2} x_iEx_{i+1} \right).$$

We will show that T' is consistent. For every natural number n, let G_n be a connected graph containing elements a_n and b_n such that the length of the shortest path from a_n to b_n is n (for example, we can take G_n to consist of a path from a_n to b_n of length n).

Now any finite subset of T' is contained in $T \cup \{\varphi_k \mid k < n\}$ for some large enough n, and $(G_n; E, a_n, b_n) \models T \cup \{\varphi_k \mid k < n\}$. So by compactness, T' is consistent.

Let $(G; E, a, b)$ be a model of T'. Since $G \models T$, G is a connected graph, but since $(G; E, a, b) \models \varphi_n$ for all n, there is no path of any length from a to b. This is a contradiction.

Example 5.9. We will show that there is no first-order formula $\varphi(x)$ in the language $\mathcal{L}_{\text{Ab}}$ of abelian groups such that for every abelian group G, $\varphi(G)$ is the torsion subgroup of G (the set of elements of finite order).

Suppose for contradiction that such a formula $\varphi(x)$ exists. Let $\mathcal{L} = \mathcal{L}_{\text{Ab}} \cup \{c\}$, where c is a new constant symbol, and let

$$T = T_{\text{Ab}} \cup \{\varphi(c)\} \cup \{nc \neq 0 \mid n \geq 1\},$$

where nc is shorthand for the term $\underbrace{c + \cdots + c}_{n \text{ times}}$.

For any finite subset $T_0 \subseteq T$, pick $m \geq 1$ such that m is greater than all n such that $nc \neq 0 \in T_0$. Then the cyclic group C_m can be expanded to a model of T_0 by interpreting c as a generator. Note that $C_m \models T_{\text{Ab}}$ and $C_m \models \varphi(c)$, since every element of C_m is torsion.

By compactness, T has a model G. Since $G \models T_{\text{Ab}}$, G is an abelian group. Since $G \models \varphi(c)$, c^G is a torsion element of G. But since $G \models nc \neq 0$ for all $n \geq 1$, c^G does not have any finite order. This is a contradiction.

5.2 The Löwenheim–Skolem theorems

Echoing Corollary 3.15, we can use compactness to build arbitrarily large elementary extensions of any infinite structure.

Theorem 5.10 (Upwards Löwenheim–Skolem)**.** *Suppose M is an infinite $\mathcal{L}$-structure and κ is a cardinal. Then there is an elementary embedding $M \to N$ with $|N| \geq \kappa$.*

Proof. Recall that $\mathcal{L}(M)$ is the language obtained by adding a new constant symbol to $\mathcal{L}$ for every element of M. Consider the language

$$\mathcal{L}' = \mathcal{L}(M) \cup \{c_\alpha \mid \alpha < \kappa\}.$$

Let T be the $\mathcal{L}'$-theory

$$\text{Diag}^{\text{FO}}(M) \cup \{c_\alpha \neq c_\beta \mid \alpha < \beta < \kappa\}.$$

A finite subset $T_0 \subseteq_{\text{fin}} T$ is contained in the theory

$$\text{Diag}^{\text{FO}}(M) \cup \{c_{\alpha_i} \neq c_{\alpha_j} \mid 1 \leq i < j \leq n\}$$

for some finitely many $\alpha_1 < \cdots < \alpha_n < \kappa$. Since M is infinite, we can pick n distinct elements $a_1, \ldots, a_n \in M$. Then interpreting the constant c_{α_i} as a_i, M itself is a model of T_0.

By compactness, T is satisfiable. Let $N \models T$. Then by Proposition 2.19, the map $a \mapsto a^N$ is an elementary embedding $M \to N$. And $|N| \geq \kappa$, since the elements $(c_\alpha^N)_{\alpha<\kappa}$ are all distinct. □

As a consequence of Theorem 5.10, if T is a theory with infinite models, it has models of arbitrarily large cardinalities. What if we want a model of cardinality exactly κ for some infinite κ? For this, we need a tool to find elementary substructures.

Definition 5.11. A substructure M of N is an **elementary substructure**, written $M \preceq N$, if the inclusion map $M \to N$ is an elementary embedding.

The following is a criterion for a substructure to be elementary. The usefulness of this criterion comes from the fact that condition (2) only refers to truth in the larger structure N – we never have to think about truth in M! But there is also a difficulty in applying this criterion, namely that it quantifies over *all* formulas $\varphi(x, y)$.

Theorem 5.12 (Tarski–Vaught Test)**.** *Suppose M is a substructure of N. The following are equivalent:*

(1) $M \preceq N$.

(2) For every formula $\varphi(x, y)$ (where x is a finite set of variables and y is a single variable) and every tuple $a \in M^x$, if $N \models \exists y\, \varphi(a, y)$, then there is some element $b \in M$ such that $N \models \varphi(a, b)$.

Proof. (1)⇒(2): Assume $N \models \exists y\, \varphi(a, y)$, with $a \in M^x$. Since $M \preceq N$, $M \models \exists y\, \varphi(a, y)$. So there is some $b \in M$ such that $M \models \varphi(a, b)$, and since $M \preceq N$, also $N \models \varphi(a, b)$.

(2)⇒(1): We prove by induction on formulas $\varphi(x)$ that for all $a \in M^x$, $M \models \varphi(a)$ if and only if $N \models \varphi(a)$.

The base case, when $\varphi(x)$ is atomic, is handled by the fact that M is a substructure of N. Then the inclusion $M \to N$ is an embedding, which preserves and reflects atomic formulas.

The inductive steps for Boolean combinations are straightforward. So we consider the case when $\varphi(x)$ is $\exists y\, \psi(x, y)$.

Assume $M \models \varphi(a)$. Then there is some $b \in M$ such that $M \models \psi(a, b)$. By induction $N \models \psi(a, b)$, so $N \models \varphi(a)$.

Conversely, assume $N \models \varphi(a)$, i.e. $N \models \exists y\, \psi(a, y)$. By (2), there is some $b \in M$ such that $N \models \psi(a, b)$. By induction, $M \models \psi(a, b)$, so $M \models \varphi(a)$. □

Theorem 5.13 (Downwards Löwenheim–Skolem)**.** *Suppose M is a structure and $A \subseteq M$. Then there is an elementary substructure $N \preceq M$ such that $A \subseteq N$ and $|N| \leq \max(\aleph_0, |A|, |\mathcal{L}|)$.*

Proof. We define a sequence $(A_i)_{i\in\omega}$ of subsets of M by recursion. Let $A_0 = A$.

Given A_i, for every formula $\varphi(x, y)$ (where x is a tuple of variables and y is a single variable) and every $a \in A_i^x$, if $M \models \exists y\, \varphi(a, y)$, pick some element $b_{\varphi(a,y)} \in M$ such that $M \models \varphi(a, b)$. Define

$$A_{i+1} = A_i \cup \{b_{\varphi(a,y)} \mid M \models \exists y\, \varphi(a, y), a \in A_i^x\}.$$

Finally, let $N = \bigcup_{i\in\omega} A_i$.

I claim that N is a substructure of M, i.e., that N is $\mathcal{L}$-closed. So suppose $f \in \mathcal{L}$ is an n-ary function symbol, and let $a \in N^n$. Since a is a finite tuple, there is some $i \in \omega$ such that $a \in A_i^n$. Then $M \models \exists y\, (f(a) = y)$, so there is some $b \in A_{i+1} \subseteq N$ such that $M \models (f(a) = b)$. Thus $f^M(a) = b \in N$.

Next, we show that $N \preceq M$ using the Tarski–Vaught test, Theorem 5.12. So suppose $\varphi(x, y)$ is a formula and $a \in N^x$, such that $M \models \exists y\, \varphi(a, y)$. Then there is some $i \in \omega$ such that $a \in A_i^x$, and by construction there is some $b \in A_{i+1} \subseteq N$ such that $M \models \varphi(a, b)$, as was to be shown.

It remains to bound the cardinality of N. Let $\kappa = \max(\aleph_0, |A|, |\mathcal{L}|)$. We show by induction that $|A_i| \leq \kappa$ for all $i \in \omega$. In the base case, $A_0 = A$, and the inequality is clear. So we consider A_{i+1}. For any formula $\varphi(x, y)$, where x is a set of n variables, the set $B_{\varphi(x,y)} = \{b_{\varphi(a,y)} \mid M \models \exists y\, \varphi(a, y), a \in A_i^x\}$

has cardinality at most $|A_i|^n$. This is equal to $|A_i|$ when A_i is infinite, and it is finite when A_i is finite. So in either case, $|B_{\varphi(x,y)}| \leq \max(|A_i|, \aleph_0) \leq \kappa$ by induction, and since κ is infinite.

Now the number of formulas $\varphi(x, y)$ is at most $\max(\aleph_0, |\mathcal{L}|) \leq \kappa$, so:

$$\begin{aligned} |A_{i+1}| &= \left| A_i \cup \bigcup_{\varphi(x,y)} B_{\varphi(x,y)} \right| \\ &\leq |A_i| + \max(\aleph_0, |\mathcal{L}|, \kappa) \\ &\leq \kappa. \end{aligned}$$

Finally, we have $|N| = |\bigcup_{i\in\omega} A_i| \leq \max(\aleph_0, \kappa) = \kappa$. □

Theorem 5.13 was the source of Skolem's "paradox". In modern language: ZFC set theory proves that there are uncountably infinite sets (like $\mathbb{R}$ and $\mathcal{P}(\omega)$). But if ZFC is consistent, then it has a countably infinite model. How can a countably infinite model of set theory contain uncountably infinite sets? The resolution of the paradox is that if M is a countable model of ZFC, then working *outside* the model, we can put M in bijection with the *real* natural numbers. But *in* M, there is no element of M which is a bijection between the element of M called $\mathcal{P}(\omega)$ and the element of M called ω.

Corollary 5.14. *Let T be a theory with at least one infinite model, and suppose $\kappa \geq \max(\aleph_0, |\mathcal{L}|)$. Then T has a model of cardinality exactly κ.*

Proof. Let M be an infinite model of T. By Theorem 5.10, there is an elementary embedding $M \to M'$ (so $M' \models T$) with $|M'| \geq \kappa$.

If $|M'| = \kappa$, we are done. Otherwise, $\max(\aleph_0, |\mathcal{L}|) \leq \kappa < |M'|$. Let A be an arbitrary subset of M' of cardinality κ. By Theorem 5.13, there exists $N \preceq M'$ (so $N \models T$) such that $A \subseteq N$ and $|N| \leq \max(\aleph_0, |A|, |\mathcal{L}|) = \kappa$, and since $A \subseteq N$, also $\kappa \leq |N|$. Thus N is model of T of cardinality exactly κ. □

Note that the Löwenheim–Skolem theorems do not guarantee existence of models of cardinality κ when $\kappa < \max(\aleph_0, |\mathcal{L}|)$. Model theory has very little to say in general about models which are finite or smaller than the cardinality of the language.

5.3 Partial types

Lecture 13: 10/27

A **partial type** (in context x) is a set of formulas in the same context x. We say that a partial type $\Sigma(x)$ is **realized** by $a \in M^x$, written $M \models \Sigma(a)$, if $M \models \varphi(a)$ for all $\varphi(x) \in \Sigma(x)$. We say that $\Sigma(x)$ is **satisfiable** relative to a theory T if it is realized in some model $M \models T$.

A **complete type** $p(x)$ relative to a theory T is a satisfiable partial type such that for every formula $\varphi(x)$ in context x, either $\varphi(x) \in p(x)$, or $\neg\varphi(x) \in p(x)$.

For any $M \models T$ and any $a \in M^x$, we define $\mathrm{tp}(a) = \{\varphi(x) \mid M \models \varphi(a)\}$. Then $\mathrm{tp}(a)$, the **complete type of** a is a complete type in context x.

Proposition 5.15 (Compactness for partial types)**.** *Let* $\Sigma(x)$ *be a partial type such that every finite subset* $\Sigma'(x) \subseteq_{\text{fin}} \Sigma(x)$ *is satisfiable relative to* T*. Then* $\Sigma(x)$ *is satisfiable relative to* T*.*

Proof. Let $\mathcal{L}' = \mathcal{L}(c)$, where c is a new set of constant symbols, one for every variable in x. Let T' be the $\mathcal{L}$-theory $T \cup \Sigma(c)$. By hypothesis, every finite subset of T' is satisfiable: For any $\Sigma'(x) \subseteq_{\text{fin}} \Sigma(x)$, if $N \models T$ and $N \models \Sigma'(a)$, then letting N' be the expansion of N in which $c^{N'} = a$, we have $N' \models T \cup \Sigma'(c)$. By compactness, there exists $M' \models T'$. Letting $a = c^{M'}$, we have $M' \models T$ and $M' \models \Sigma(a)$. □

We also want to consider types with parameters. Given a set $A \subseteq M$, we can view a formula like $\varphi(x, b)$, where b is a tuple from A, as an $\mathcal{L}(A)$-formula. A **partial type over** A is just a partial type in the language $\mathcal{L}(A)$.

Theorem 5.16. *Suppose* $A \subseteq M$ *and* $\Sigma(x)$ *is a partial type over* A*. Then* $\Sigma(x)$ *is finitely satisfiable in* M *if and only if it is realized in some elementary extension of* M*.*

Proof. Suppose $\Sigma(x)$ is finitely satisfiable in M. Then for every finite $\Sigma'(x) \subseteq \Sigma(x)$, $\Sigma'(x)$ is realized by some $a \in M^x$. Then $\Sigma(x)$ is finitely satisfiable relative to $\text{Diag}^{\text{FO}}(M)$, so by Proposition 5.15, $\Sigma(x)$ is satisfiable relative to $\text{Diag}^{\text{FO}}(M)$. Let $N \models \text{Diag}^{\text{FO}}(M)$ and $a \in N^x$ realize Σ. Since $N \models \text{Diag}^{\text{FO}}(M)$, there is an elementary embedding $M \to N$, and identifying M with its image in N, Σ is realized in an elementary extension of M.

Conversely, suppose $\Sigma(x)$ is realized by $a \in N^x$ with $M \preceq N$. Let $\Sigma' \subseteq_{\text{fin}} \Sigma$, and let $\sigma = \bigwedge_{\psi \in \Sigma'} \psi$. Then since $N \models \sigma(a)$, $N \models \exists x\, \sigma(x)$. Since σ is an $\mathcal{L}(A)$-formula and $A \subseteq M \preceq N$, we have $M \models \exists x\, \sigma(x)$. Letting $b \in M^x$ with $M \models \sigma(b)$, b realizes $\Sigma'(x)$ in M. So Σ is finitely satisfiable in M. □

Definition 5.17. Let T be a theory. Partial types $\Sigma(x)$ and $\Sigma'(x)$ in the same context x are T-equivalent if for all $M \models T$ and all $a \in M^x$, $M \models \Sigma(a)$ if and only if $M \models \Sigma'(a)$. We extend the definition in the obvious way to formulas, treating a formula $\varphi(x)$ as the partial type $\{\varphi(x)\}$.

Lemma 5.18. *Suppose* $\Sigma(x)$ *is a partial type which is* T*-equivalent to a formula* $\varphi(x)$*. Then there is a finite subset* $\Sigma' \subseteq \Sigma$ *such that* Σ' *is* T*-equivalent to* φ*.*

Proof. Since Σ and φ are T-equivalent, the partial type $\Sigma \cup \{\neg\varphi\}$ is not satisfiable relative to T. By Proposition 5.15, there is a finite $\Sigma' \subseteq_{\text{fin}} \Sigma$ such that $\Sigma' \cup \{\neg\varphi\}$ is not satisfiable relative to T. I claim that Σ' and φ are T-equivalent.

Let $M \models T$ and $a \in M^x$. If $M \models \varphi(a)$, then since φ and Σ are equivalent, $M \models \Sigma(a)$, so $M \models \Sigma'(a)$. Conversely, if $M \models \Sigma'(a)$, then since $\Sigma' \cup \{\neg\varphi\}$ is not satisfiable relative to T, we must have $M \models \varphi(a)$. □

5.4 Preservation results

We have already seen (Theorem 2.13) that $\mathcal{F}$-morphisms preserve all $\exists^+(\mathcal{F})$-formulas and (Theorem 4.6) that direct limits along $\mathcal{F}$-morphisms preserve

$\exists^+(\mathcal{F})$-theories. Our goal in this section is to establish the converses to these facts: if a formula is preserved by $\mathcal{F}$-morphisms, it is equivalent to an $\exists^+(\mathcal{F})$-formula, and if a theory is preserved by direct limits along $\mathcal{F}$-morphisms, then it is equivalent to an $\exists^+(\mathcal{F})$-theory.

We begin with a very useful amalgamation result. Recall that $h\colon M \dashrightarrow N$ is a partial $\mathcal{F}$-morphism if the domain of h is $A \subseteq M$ and for all $\mathcal{F}$-formulas $\varphi(x)$ and $a \in A^x$, if $M \models \varphi(a)$, then $N \models \varphi(h(a))$.

Theorem 5.19. *Let $h\colon M \dashrightarrow N$ be a partial $\exists^+(\mathcal{F})$-morphism. Then there is an elementary extension $N \preceq N'$ and a $\mathcal{F}$-morphism $g\colon M \to N'$ extending h.*

Proof. Let $A \subseteq M$ be the domain of h. Consider the following $\mathcal{L}(M \cup N)$-theory T:

$$\mathrm{Diag}^{\mathcal{F}}(M) \cup \mathrm{Diag}^{\mathrm{FO}}(N) \cup \{a = n \mid a \in A, h(a) = n\}.$$

It suffices to show that T is satisfiable. Indeed, if $N' \models T$, then because $N' \models \mathrm{Diag}^{\mathrm{FO}}(N)$, up to replacing N' with an isomorphic model, N' as an elementary extension of N. And since $N' \models \mathrm{Diag}^{\mathcal{F}}(M)$, there is an $\mathcal{F}$-morphism $g\colon M \to N'$, defined by $g(m) = m^{N'}$. Finally, for $a \in A$, if $h(a) = n$, then $N \models a = n$, so $g(a) = a^{N'} = n^{N'} = n = h(a)$, and thus g extends h.

Let $T_0 \subseteq T$ be a finite subtheory. Let a be the finite set of elements of A appearing in T_0, and let m be the finite set of elements of $M \setminus A$ appearing in T_0. Let $\psi(a, m)$ be the conjunction of the finitely many formulas from $\mathrm{Diag}^{\mathcal{F}}(M)$ in T_0. Then $M \models \psi(a, m)$, so $M \models \exists y\, \psi(a, y)$. Since h is an $\exists^+(\mathcal{F})$-morphism, $N \models \exists y\, \psi(h(a), y)$. Let $n \in N^y$ be such that $N \models \psi(h(a), n)$. Then interpreting $a^N = h(a)$ and $m^N = n$, $N \models T_0$. By compactness, T is satisfiable. □

Lecture 14: 10/29

Corollary 5.20. *Let $f\colon M \to N$ and $f'\colon M \to N'$ be elementary embeddings. Then there exist elementary embeddings $g\colon N \to M'$ and $g'\colon N' \to M'$ such that $g \circ f = g' \circ f'$.*

Proof. Let $h\colon N \dashrightarrow N'$ be the partial map with domain $f(M) \subseteq N$, defined by $h(f(m)) = f'(m)$. Then h is a partial FO-morphism. Indeed, for all formulas $\varphi(x)$ and $a \in f(M)^x$, writing $a = f(m)$, if $N \models \varphi(a)$, then $N \models \varphi(f(m))$, so $M \models \varphi(m)$, and $N' \models \varphi(f'(m))$, so $N' \models \varphi(h(a))$. Since $\exists^+(\mathrm{FO}) = \mathrm{FO}$, by Theorem 5.19, there exists an elementary extension $N' \preceq M'$ and an elementary embedding $g\colon N \to M'$ extending h. Letting g' be the elementary inclusion map $N' \to M'$, this means that for all $m \in M$, $g(f(m)) = h(f(m)) = f'(m) = g'(f'(m))$, so $g \circ f = g' \circ f'$. □

Theorem 5.21. *Let T be a theory and $\Sigma(x)$ a partial type which is preserved by all $\mathcal{F}$-morphisms between models of T. Then $\Sigma(x)$ is T-equivalent to a set of $\exists^+(\mathcal{F})$-formulas.*

Proof. Let $\Sigma'(x)$ be the set of all $\exists^+(\mathcal{F})$-formulas $\varphi(x)$ such that $T \cup \Sigma(x) \models \varphi(x)$. Then if $M \models T$ and $a \in M^x$ realizes $\Sigma(x)$, we have $M \models \varphi(a)$ for all $\varphi \in \Sigma'$, so a realizes $\Sigma'(x)$.

Conversely, assume $M \models T$ and $a \in M^x$ realizes $\Sigma'(x)$. Let

$$\Delta(x) = \{\neg\varphi(x) \mid \varphi \in \exists^+(\mathcal{F}) \text{ and } M \not\models \varphi(a)\}.$$

I claim that $\Sigma(x) \cup \Delta(x)$ is satisfiable relative to T. If not, then by compactness, there is a finite subset $\Delta_0 \subseteq \Delta$ such that $\Sigma \cup \Delta_0$ is not satisfiable relative to T. Let $\Delta_0 = \{\neg\varphi_i \mid 1 \leq i \leq n\}$. Then $T \cup \Sigma(x) \models \bigvee_{i=1}^n \varphi_i(x)$, so $\bigvee_{i=1}^n \varphi_i(x) \in \Sigma'(x)$. This contradicts the fact that $M \not\models \varphi_i(a)$ for all $1 \leq i \leq n$.

Now let $N \models T$ and $b \in N^x$ realizing $\Sigma(x) \cup \Delta(x)$. If $\chi(x)$ is an $\exists^+(\mathcal{F})$-formula and $N \models \chi(b)$, then $\neg\chi(b) \notin \Delta$, so $M \models \chi(a)$. Thus the partial function $h\colon N \dashrightarrow M$ mapping b to a is a partial $\exists^+(\mathcal{F})$-morphism. By Theorem 5.19, there is an elementary extension $M \preceq M'$ and an $\mathcal{F}$-morphism $g\colon N \to M'$ extending g. In particular, $g(b) = a$. Since $N \models \Sigma(b)$ and Σ is preserved by $\mathcal{F}$-morphisms, $M' \models \Sigma(a)$, and since $M \preceq M'$, $M \models \Sigma(a)$, as desired. □

Corollary 5.22. *Let T' be a theory which is preserved by all $\mathcal{F}$-morphisms. Then T' is equivalent to a set of $\exists^+(\mathcal{F})$-sentences.*

Proof. Apply Theorem 5.21 with $T = \varnothing$ to T', viewed as a partial type in the empty context. □

Corollary 5.23. *Let $\varphi(x)$ be a formula which is preserved by all $\mathcal{F}$-morphisms between models of T. Then $\varphi(x)$ is T-equivalent to an $\exists^+(\mathcal{F})$-formula.*

Proof. Apply Theorem 5.21 to the partial type $\{\varphi(x)\}$. We obtain that $\varphi(x)$ is T-equivalent to a set $\Sigma(x)$ of $\exists^+(\mathcal{F})$-formulas. Now by Lemma 5.18, since $\Sigma(x)$ is T-equivalent to a formula $\varphi(x)$, there is a finite $\Sigma' \subseteq_{\text{fin}} \Sigma$ which is equivalent to $\varphi(x)$. Then $\varphi(x)$ is equivalent to the $\exists^+(\mathcal{F})$-formula $\bigwedge_{\sigma \in \Sigma'} \sigma$. □

Definition 5.24. An **$\mathcal{F}$-immersion** is an $\mathcal{F}$-morphism which also reflects $\exists^+(\mathcal{F})$-formulas.

Let $\text{Diag}^{\neg\exists^+(\mathcal{F})}(M) = \{\neg\varphi \mid \varphi \notin \text{Diag}^{\exists^+(\mathcal{F})}(M)\}$. Since reflecting all $\exists^+(\mathcal{F})$-formulas is the same as preserving their negations, we can build an $\mathcal{F}$-immersion out of M by finding a model of $\text{Diag}^{\mathcal{F}}(M) \cup \text{Diag}^{\neg\exists^+(\mathcal{F})}(M)$.

Theorem 5.25. *Let T be a theory which is preserved by direct limits along $\mathcal{F}$-morphisms, in the sense that if $(M_i)_{i \in I}$ is a directed family of models of T with all connecting maps f_{ij} are $\mathcal{F}$-morphisms, then $\varinjlim M_i \models T$. Then T is equivalent to an $\exists^+(\mathcal{F})$-theory (i.e., a set of $\exists^+(\mathcal{F})$-sequents).*

Proof. Let T' be the set of all $\exists^+(\mathcal{F})$-sequents S such that $T \models S$. Then every model of T is a model of T'. We must show the converse, that every model of T' is a model of T.

Let $M \models T'$. We build a diagram of the following shape:

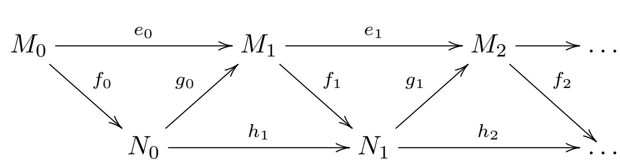

such that:

- $M_0 = M$ and each $M_i \models T'$.
- Each $N_i \models T$.
- Each e_i is an elementary embedding.
- Each f_i is an $\mathcal{F}$-immersion.
- Each g_i and each h_i is an $\mathcal{F}$-morphism.
- The triangles commute: $g_i \circ f_i = e_i$ and $f_{i+1} \circ g_i = h_{i+1}$ for all i.

In the base case, we let $M_0 = M \models T'$.

Given $M_i \models T'$, we build $f_i \colon M_i \to N_i \models T$. Consider the $\mathcal{L}(M_i)$-theory:

$$T \cup \mathrm{Diag}^{\mathcal{F}}(M_i) \cup \mathrm{Diag}^{\neg\exists^+(\mathcal{F})}(M_i)$$

If this theory is inconsistent, then there are finitely many $\varphi_j \in \mathrm{Diag}^{\mathcal{F}}(M_i)$ for $1 \le j \le n$ and finitely many $\neg\psi_k \in \mathrm{Diag}^{\neg\exists^+(\mathcal{F})}(M_i)$ for $1 \le k \le m$ such that

$$T \cup \left\{ \bigwedge_{j=1}^{n} \varphi_j, \bigwedge_{k=1}^{m} \neg\psi_k \right\}$$

is inconsistent. Let c be the finite set of all constant symbols not in $\mathcal{L}$ that are mentioned in the φ_j and ψ_k, and let z be a set of variables corresponding to the constant symbols in c. For any $A \models T$ and $a \in A^z$, if $A \models \bigwedge_{j=1}^{n} \varphi_i(a)$, then $A \not\models \bigwedge_{k=1}^{m} \neg\psi_k(a)$, so

$$T \models \left(\bigwedge_{j=1}^{n} \varphi_j(z) \vdash_z \bigvee_{k=1}^{m} \psi_k(z) \right).$$

This is an $\exists^+(\mathcal{F})$-sequent, so it is in T'. Since $M_i \models T'$, this contradicts the fact that $M \models \bigwedge_{j=1}^{n} \varphi_j(c) \wedge \neg \bigvee_{k=1}^{m} \psi_k(c)$.

Letting N_i be a model for the theory above, the map $f_i(m) = m^{N_i}$ is an $\mathcal{F}$-immersion, as desired.

Given M_i and N_i and the $\mathcal{F}$-immersion $f_i \colon M_i \to N_i$, we build an elementary extension $M_i \preceq M_{i+1}$ and a $\mathcal{F}$-morphism $g_i \colon N_i \to M_{i+1}$. Note that since f_i is injective (since it reflects the $\mathcal{F}$-formula $x = y$), it is a bijection onto its range. Thus we can view f_i^{-1} as a partial map $N_i \dashrightarrow M_i$. Since f_i is a $\mathcal{F}$-immersion, f_i^{-1} is a partial $\exists^+(\mathcal{F})$-morphism. By Theorem 5.19, f_i^{-1} extends to an $\mathcal{F}$-morphism $g_i \colon N_i \to M_{i+1}$ for some elementary extension $M_i \preceq M_{i+1}$.

Now we define e_i to be the elementary inclusion map $M_i \to M_{i+1}$. The fact that g_i extends f_i^{-1} implies that for all $m \in M_i$, $g_i(f_i(m)) = f_i^{-1}(f_i(m)) = m = e_i(m)$, so $e_i = g_i \circ f_i$.

Finally, for $i > 0$, we define $h_i = f_i \circ g_{i-1}$. Since both f_i and g_{i-1} are $\mathcal{F}$-morphisms, so is h_i.

Having build the diagram, let $M = \varinjlim M_i$ and $N = \varinjlim N_i$. The homomorphisms f_i induce a homomorphism $M \to N$, and the homomorphisms g_i induce a homomorphism $N \to M$, which are inverses. Thus $M \cong N$.

Since $N_i \models T$ for all i, each h_i is an $\mathcal{F}$-morphism, and T is preserved by direct limits along $\mathcal{F}$-morphisms, $N \models T$. Since $M \cong N$, $M \models T$. And since M is a direct limit of the M_i along elementary embeddings, the canonical map $M_0 \to M$ is an elementary embedding. Thus $M_0 \models T$, as desired. □

6 Model completeness

Lecture 15: 11/3

In this section, we apply the results of the previous sections with $\mathcal{F} = \mathrm{Lit}$, the literal fragment, consisting of atomic and negated atomic formulas. Thus, $\mathcal{F}$-morphisms are embeddings, and $\exists^+(\mathcal{F})$ is the existential fragment $\exists$, built from literals by $\exists$, $\wedge$, and $\vee$, immersions are embeddings that reflect $\exists$-formulas, etc.

By Theorem 5.25, a first-order theory T is preserved by direct limits along embeddings if and only if T is equivalent to a set of $\exists$-sequents. By syntactic manipulations, any set of $\exists$-sequents is equivalent to a set of $\forall\exists$-sentences (sentences in the closure of the $\exists$-fragment under $\forall$, $\wedge$, and $\vee$) and vice versa. To avoid any ambiguity between $\exists$-theories (in the sense of a set of $\exists$-sequents) and $\forall\exists$-theories (in the sense of a set of $\forall\exists$-sentences), we refer to such theories as **inductive theories** (since their class of models are closed under *inductive limits*, an alternate term for direct limits).

6.1 Characterizing model complete theories

One difficulty in doing model theory with arbitrary first-order theories is that many concepts quantify over all formulas, and arbitrary first-order formulas (with arbitrary quantifier-alternations) can be very complicated. For example, to qualify as an elementary embedding, a function must preserve a lot of structure! It is much more convenient to work with model-complete theories, where this complexity is reduced to a level that is much more manageable to check, both for elementary embeddings and for definability.

Definition 6.1. A theory T is **model complete** if every embedding between models of T is an elementary embedding.

Definition 6.2. Let $\mathcal{K}$ be a class of $\mathcal{L}$-structures (usually the class of models of a theory T). We say that a structure $M \in \mathcal{K}$ is **existentially closed** (or EC) in $\mathcal{K}$, if for every embedding $f\colon M \to N$ with $N \in \mathcal{K}$ is an immersion (reflects $\exists$-formulas).

Condition (6) in the following theorem is the reason for the name "model-complete".

Theorem 6.3. *Let T be a theory. The following are equivalent:*

(1) T is model complete.

(2) Every model of T is EC.

(3) Every $\exists$-formula is T-equivalent to a $\forall$-formula.

(4) Every formula is T-equivalent to a $\forall$-formula.

(5) Every formula is T-equivalent to an $\exists$-formula.

(6) For every model $M \models T$, the $\mathcal{L}(M)$-theory $T \cup \mathrm{Diag}(M)$ is a complete $\mathcal{L}(M)$-theory.

Proof. (1)⇒(2): Suppose $M \models T$ and $h\colon M \to N$ is an embedding with $N \models T$. Let φ be an $\exists$-formula. Since h is an elementary embedding by (1), h reflects φ. Thus h is an immersion and M is an EC model of T.

(2)⇒(3): Let φ be an $\exists$-formula. Every embedding between models of T reflects φ, and hence preserves $\neg\varphi$. By Corollary 5.23, $\neg\varphi$ is T-equivalent to an $\exists$-formula, so φ is T-equivalent to a $\forall$-formula.

(3)⇒(4): Note since every $\exists$-formula is T-equivalent to a $\forall$-formula by (3), taking negations, every $\forall$-formula is T-equivalent to an $\exists$-formula. Let's call this statement (3').

We prove by induction that every formula is T-equivalent to a $\forall$-formula. The cases of atomic formulas and $\top$ and $\bot$ are trivial, since these are $\forall$-formulas. The inductive steps for $\forall$, $\wedge$, and $\vee$ are trivial, since the $\forall$-fragment is closed under these operations.

Suppose $\varphi(x)$ is $\neg\psi(x)$. By induction, we may assume that $\psi(x)$ is a $\forall$-formula. By (3'), $\psi(x)$ is T-equivalent to an $\exists$-formula $\psi'(x)$, so $\varphi(x)$ is T-equivalent to $\neg\psi'(x)$, which is equivalent to a $\forall$-formula.

Suppose $\varphi(x)$ is $\exists y\,\psi(x,y)$. By induction, we may assume that $\psi(x,y)$ is a $\forall$-formula. By (3'), $\psi(x,y)$ is T-equivalent to an $\exists$-formula $\psi'(x,y)$, so $\varphi(x)$ is T-equivalent to $\exists y\,\psi'(x,y)$, which is an $\exists$-formula. By (3), $\varphi(x)$ is T-equivalent to a $\forall$-formula.

(4)⇒(5): For any formula φ, $\neg\varphi$ is equivalent to a $\forall$-formula ψ by (4), so φ is equivalent to $\neg\psi$, which is equivalent to an $\exists$-formula.

(5)⇒(1): Let $h\colon M \to N$ be an embedding between models of T. Let φ be an arbitrary formula. Since φ is T-equivalent to an $\exists$-formula by (5), φ is preserved by h. So h is an elementary embedding.

(1)⇒(6): Let $M \models T$. We would like to show that $T \cup \mathrm{Diag}(M) \models \mathrm{Diag}^{\mathrm{FO}}(M)$, since $\mathrm{Diag}^{\mathrm{FO}}(M)$ is a complete $\mathcal{L}(M)$-theory. So let $\varphi(m) \in \mathrm{Diag}^{\mathrm{FO}}(M)$ (with the constant symbols m naming elements of M made explicit). Let $N \models T \cup \mathrm{Diag}(M)$. Then the natural embedding $h\colon M \to N$ is an elementary embedding by (1), so it preserves $\varphi(x)$, and we have $N \models \varphi(m)$. It follows that $T \cup \mathrm{Diag}(M) \models \varphi(m)$.

(6)⇒(1): Let $h\colon M \to N$ be an embedding between models of T. Then $N \models T \cup \mathrm{Diag}(M)$, where we interpret the constant symbols naming elements of M as their images in N under h. Let $\varphi(x)$ be an arbitrary formula and $m \in M^x$, and suppose $M \models \varphi(m)$. Then since $T \cup \mathrm{Diag}(M)$ is complete and M is a model, $T \cup \mathrm{Diag}(M) \models \varphi(m)$. It follows that $N \models \varphi(h(m))$. Thus h is an elementary embedding. □

Theorem 6.4. *Let T be a model complete theory. Then T is equivalent to an inductive theory (hence has an axiomatization by $\exists$-sequents or by $\forall\exists$-sentences).*

Proof. By Theorem 5.25, it suffices to show that the class of models of T is closed under direct limits along embeddings. Let $(M_i)_{i\in I}$ be a directed family such that $M_i \models T$ for all $i \in I$ and all connecting maps are embeddings. Since T is model complete, all connecting maps are elementary embeddings. But then, picking any $j \in I$, the map $g_j\colon M_j \to \varinjlim M_i$ is an elementary embedding, so since $M_j \models T$, also $\varinjlim M_i \models T$. □

6.2 Quantifier elimination

How can we give examples of model complete theories? If we're lucky, we might be able to prove that a theory of interest eliminates quantifiers.

Definition 6.5. A theory T has **quantifier elimination** (or **eliminates quantifiers**) if every formula is T-equivalent to a quantifier-free formula.

Example 6.6. Let T be the theory of fields. The formula $\exists y\,(xy = 1)$, expressing that x is invertible, is T-equivalent to the quantifier-free formula $x \neq 0$.

Similarly, the formula

$$\exists a\,\exists b\,\exists c\,\exists d\,(ax + bz = 1 \wedge ay + bw = 0 \wedge cx + dz = 0 \wedge cy + dw = 0),$$

expressing that the matrix $\begin{pmatrix} x & y \\ z & w \end{pmatrix}$ is invertible, is T-equivalent to the quantifier-free formula $xw - yz \neq 0$.

The formula $\varphi(a, b, c)$ defined by $\exists y\,(ay^2 + by + c = 0)$ is not T-equivalent to a quantifier-free formula. To see this, note for example that if $h\colon \mathbb{R} \to \mathbb{C}$ is the inclusion embedding, h preserves and reflects all quantifier-free formulas. But it does not reflect φ, since $\mathbb{R} \not\models \varphi(1, 0, 1)$, while $\mathbb{C} \models \varphi(1, 0, 1)$.

However, if we include $\leq$ in the language and move to the complete theory of $\mathbb{R}$, we have that φ is $\mathrm{Th}(\mathbb{R}; 0, 1, +, -, \cdot, \leq)$-equivalent to the quantifier-free formula $b^2 - 4ac \geq 0$.

We begin with a useful structural criterion for a single formula to be equivalent to a quantifier-free formula. Note that quantifier-free formulas are both $\exists$-formulas and $\forall$-formulas, so they are both preserved and reflected by embeddings.

Theorem 6.7. *Let T be a theory and $\varphi(x)$ a formula. The following are equivalent:*

(1) φ is T-equivalent to a quantifier-free formula.

(2) Suppose M and N are models of T, A is any structure, and $g\colon A \to M$ and $h\colon A \to N$ are embeddings. For all $a \in A^x$, if $N \models \varphi(h(a))$, then $M \models \varphi(g(a))$.

(3) The same as (2), but with the extra requirement that a is a finite set of generators for A.

Proof. (1)⇒(2): Suppose φ is T-equivalent to a quantifier-free formula ψ. Then in the context of (2), since g and h preserve and reflect ψ, we have

$$\begin{aligned} N \models \varphi(h(a)) &\Rightarrow N \models \psi(h(a)) \\ &\Rightarrow A \models \psi(a) \\ &\Rightarrow M \models \psi(g(a)) \\ &\Rightarrow M \models \varphi(g(a)) \end{aligned}$$

(2)⇒(3): Trivial.

(3)⇒(1): Let Ψ be the set of all quantifier-free formulas $\psi(x)$ such that $T \cup \{\varphi(x)\} \models \psi(x)$. Clearly $T \cup \{\varphi(x)\} \models \Psi(x)$. We will show that $T \cup \Psi(x) \models \{\varphi(x)\}$.

Let $M \models T$, and let $a \in M^x$ be a realization of $\Psi(x)$. Let $A = \langle a \rangle_M$, and let $g \colon A \to M$ be the inclusion. We would like to find a model for the theory $T \cup \mathrm{Diag}_M(a) \cup \{\varphi(a)\}$.

Suppose for contradiction that this theory is not satisfiable. Then by compactness, there is a finite subset which is unsatisfiable, so there is a finite subset $\Theta \subseteq_{\text{fin}} \mathrm{Diag}_M(a)$ such that $T \cup \Theta \cup \{\varphi(a)\}$ is not satisfiable. But then $T \cup \{\varphi(a)\} \models \neg \bigwedge_{\theta(a)\in\Theta} \theta(a)$, so $T \cup \{\varphi(x)\} \models \bigvee_{\theta(a)\in\theta} \neg\theta(x)$, and thus $\bigvee_{\theta(a)\in\theta} \neg\theta(x) \in \Psi$. So $M \models \bigvee_{\theta(a)\in\theta} \neg\theta(a)$, contradicting the fact that $\theta(a) \in \mathrm{Diag}_M(a)$ for all $\theta(a) \in \Theta$.

Thus there is a model $N \models T \cup \mathrm{Diag}_M(a) \cup \{\varphi(a)\}$. It follows that the map $a \mapsto a^N$ extends to an embedding $h \colon A \to N$. Now since $N \models \varphi(h(a))$, by (3), $M \models \varphi(g(a))$, and $g(a) = a$.

We have shown that the partial types $\{\varphi(x)\}$ and $\Psi(x)$ are T-equivalent. By Lemma 5.18, there is a finite subset $\Psi_0 \subseteq_{\text{fin}} \Psi$ such that $\{\varphi\}$ is T-equivalent to Ψ_0. Then $\varphi(x)$ is T-equivalent to the quantifier-free formula $\bigwedge_{\psi\in\Psi_0} \psi$. □

Note that the proof of Theorem 6.7 used the compactness theorem twice (once by the application of Lemma 5.18. Compactness is a fundamentally non-constructive proof technique, so the proof doesn't give us any information on how to explicitly find the quantifier-free formula equivalent to φ. In certain circumstances, we can do better by giving an explicit algorithm for eliminating quantifiers. Lecture 16: 11/5

Next we note that to prove that T has quantifier elimination, we only have to consider formulas of a particularly simple form. This is the analogue of the result that a theory is model complete if and only if every $\exists$-formula is equivalent to a $\forall$-formula.

Definition 6.8. A **primitive formula** is one of the form $\exists y \bigwedge_{i=1}^n \varphi_i$, where y is a single variable and each formula φ_i is literal.

Theorem 6.9. *If every primitive formula is T-equivalent to a quantifier-free formula, then T has quantifier elimination.*

Proof. We prove by induction on the complexity of formulas that every formula is T-equivalent to a quantifier-free formula. The cases of atomic formulas, $\top$, $\bot$, $\wedge$, $\vee$, and $\neg$ are clear, and a formula of the form $\forall y\, \varphi$ can be rewritten as $\neg\exists y\, \neg\varphi$, so it suffices to handle the induction step for the existential quantifier.

So consider a formula φ of the form $\exists y\, \psi$. By induction, ψ is T-equivalent to a quantifier-free formula θ. Writing θ in disjunctive normal form, φ is T-equivalent to

$$\exists y \left(\bigvee_{i=1}^{n} \bigwedge_{j=1}^{m} \varphi_{ij} \right),$$

where each φ_{ij} is literal. It follows that φ is T-equivalent to

$$\bigvee_{i=1}^{n} \exists y \left(\bigwedge_{j=1}^{m} \varphi_{ij} \right),$$

since existential quantifiers distribute over disjunctions. For fixed i, the formula $\exists y \left(\bigwedge_{j=1}^{m} \varphi_{ij} \right)$ is primitive, so by hypothesis it is T-equivalent to a quantifier-free formula, and φ is T-equivalent to the disjunction of these n formulas. $\square$

We can now give several tests for quantifier elimination.

Theorem 6.10. *Let T be a theory. The following are equivalent:*

(1) T has quantifier elimination.

*(2) T is **substructure complete**: For any $M \models T$ and $A \subseteq M$ a substructure, the $\mathcal{L}(A)$-theory $T \cup \mathrm{Diag}(A)$ is complete.*

(3) Let $M, N \models T$ and $A \subseteq M$ a substructure. Let $f\colon A \to N$ be an embedding. Then there exists an elementary extension $N \preceq N'$ and an embedding $g\colon M \to N'$ extending f.

(4) Let $M, N \models T$ and $A \subseteq M$ a finitely generated substructure. Let $f\colon A \to N$ be an embedding. Let $b \in M$, and let $B = \langle A \cup \{b\} \rangle_M$. Then there exists an elementary extension $N \preceq N'$ and an embedding $g\colon B \to N'$ extending f.

(5) Suppose M and N are models of T, $\varphi(x)$ is a primitive formula, and $a \in M^x$ such that $M \models \varphi(a)$. Let $A = \langle a \rangle_M$, and let $f\colon A \to N$ be an embedding. Then $N \models \varphi(f(a))$.

Proof. (1)⇒(2): Assume T has quantifier elimination, $M \models T$, and $A \subseteq M$. To show that $T \cup \mathrm{Diag}(A)$ is complete, it suffices to show that $T \cup \mathrm{Diag}(A) \models \mathrm{Diag}^{\mathrm{FO}}(A)$, since the latter is a complete $\mathcal{L}(A)$-theory. Let $\varphi(a) \in \mathrm{Diag}^{\mathrm{FO}}_M(A)$, and let $N \models T \cup \mathrm{Diag}(A)$. We need to show $N \models \varphi(a)$.

By quantifier elimination, $\varphi(x)$ is T-equivalent to a quantifier-free formula $\psi(x)$. Since $M \models T$ and $M \models \varphi(a)$, we have $M \models \psi(a)$, and since ψ is quantifier-free, $A \models \psi(a)$.

Write $\psi(x)$ in disjunctive normal form as $\bigvee_{i=1}^{m} \bigwedge_{j=1}^{n} \psi_{ij}(x)$. Then there is some $1 \leq i \leq m$ such that $A \models \bigwedge_{i=1}^{n} \psi_{ij}(a)$, so $\{\psi_{ij}(a) \mid 1 \leq j \leq n\} \subseteq \mathrm{Diag}(A)$. Thus $N \models \psi_{ij}(a)$ for all $1 \leq j \leq n$, so $N \models \psi(a)$. Since $N \models T$, also $N \models \varphi(a)$, as desired.

(2)⇒(3): Suppose T is substructure complete. Let $M, N \models T$, let $A \subseteq M$ be a substructure, and let $f\colon A \to N$ be an embedding. Viewing f as a partial map $M \dashrightarrow N$, I claim that f is a partial elementary map. Let $\varphi(x)$ be a formula and $a \in A^x$, and suppose $M \models \varphi(a)$. Since $T \cup \mathrm{Diag}(A)$ is a complete $\mathcal{L}(A)$-theory, $T \cup \mathrm{Diag}(A) \vdash \varphi(a)$. Interpreting $b^n = f(b)$ for all $b \in A$, $N \models T \cup \mathrm{Diag}(A)$, so $N \models \varphi(f(a))$.

In particular, f is a partial $\exists$-morphism, so by Theorem 5.19, there exists $N \preceq N'$ and an embedding $g\colon M \to N'$ extending f.

(3)⇒(4): Trivial, by restricting the g provided by (3) from M to $B \subseteq M$.

(4)⇒(5): Let $M, N \models T$, $\varphi(x)$ a primitive formula, and $a \in M^x$ such that $M \models \varphi(a)$. Let $A = \langle a \rangle_M$, and let $f\colon A \to N$ be an embedding.

Write φ as $\exists y\, \psi(x, y)$, where ψ is a conjunction of atomic formulas. Let $b \in M$ such that $M \models \psi(a, b)$. Let $B = \langle A \cup \{b\} \rangle_M$. By (4), there exists $N \preceq N'$ and an embedding $g\colon B \to N'$ extending f. Since ψ is quantifier-free, $B \models \psi(a, b)$, and since g is an embedding, $N' \models \psi(g(a), g(b))$. Then $N' \models \exists y\, \psi(g(a), y)$, and since $N \preceq N'$, $N \models \varphi(g(a))$. Since g extends f, $g(a) = f(a)$, and we are done.

(5)⇒(1): By Theorems 6.7 and 6.9, it suffices to show that if $\varphi(x)$ is a primitive formula, $M, N \models T$, $A = \langle a \rangle$ with $a \in A^x$, $g\colon A \to M$ and $h\colon A \to N$ embeddings, and $N \models \varphi(h(a))$, then $M \models \varphi(g(a))$.

Let $A' = h(A) = \langle h(a) \rangle_N$. Let $f = g \circ h^{-1}\colon A' \to M$. Then f is an embedding. By (5), since $N \models \varphi(h(a))$, $M \models \varphi(g(a))$, as desired. □

Example 6.11. Let $\mathcal{L} = \varnothing$, and let T_{Inf} be the theory of an infinite set. We show that T_{Inf} has quantifier elimination using condition (3) from Theorem 6.10.

Let M and N be infinite sets and $A \subseteq M$. Let $f\colon A \to N$ be an embedding (an injective function). By Löwenheim–Skolem, there exists $N \preceq N'$ with $|N'| > |M|$. In particular, $|N' \setminus f(A)| = |N'| > |M| \geq |M \setminus A|$. Picking an injective function $(M \setminus A) \to (N' \setminus f(A))$, we can extend f to an injective function $g\colon M \to N'$. By Theorem 6.10, T_{Inf} has quantifier elimination.

On the other hand, the theory of arbitrary sets (the empty $\mathcal{L}$-theory) fails to have quantifier elimination – it is not even model complete. For example, the inclusion embedding $\{0\} \to \{0, 1\}$ is not elementary, because it fails to preserve the sentence $\forall x \forall y\, (x = y)$.

Example 6.12. Let $\mathcal{L} = \{<\}$, and let DLO be the theory of non-empty dense linear orders without endpoints. We show that DLO has quantifier elimination using condition (4) from Theorem 6.10.

Let M and N be models of DLO and $A \subseteq M$ a finitely generated substructure. Then A is finite, and we can enumerate its elements in increasing order: $a_1 < a_2 < \cdots < a_n$. Let $f\colon A \to N$ be an embedding. Then $f(a_1) < f(a_2) < \cdots < f(a_n)$.

Let $b \in M$, and let $B = \langle A \cup \{b\} \rangle_M = A \cup \{b\}$. We will extend f to an embedding $g\colon B \to N$. If $b \in A$, then we can take $g = f$. So assume $b \notin A$.

Case 1: If $b < a_1$, then since N has no least element, there exists $b' \in N$ with $b' < f(a_1)$. Define $g(b) = b'$.

Case 2: If $b > a_n$, then since N has no greatest element, there exists $b' \in N$ with $b' > f(a_n)$. Define $g(b) = b'$.

Case 3: Otherwise, there is some $1 \leq i < n$ such that $a_i < b < a_{i+1}$. Since $f(a_i) < f(a_{i+1})$ and N is dense, there exists $b' \in N$ with $f(a_i) < b' < f(a_{i+1})$. Define $g(b) = b'$.

By Theorem 6.10, DLO has quantifier elimination.

We will use condition (5) from Theorem 6.10 to prove that the theory of algebraically closed fields has quantifier elimination in the next section.

Proposition 6.13. *Suppose T is a satisfiable theory with quantifier elimination. The following are equivalent:*

1. *T is complete.*
2. *For any $M, N \models T$, $\langle \varnothing \rangle_M \cong \langle \varnothing \rangle_N$.*
3. *There is a structure A such that for any model $M \models T$, there is an embedding $g\colon A \to M$.*

Proof. (1)⇒(2): Let $M, N \models T$. There is a unique map $f\colon \varnothing \to \langle \varnothing \rangle_N$. Since T is complete and M and N satisfy the atomic sentences, f extends to a unique homomorphism $g\colon \langle \varnothing \rangle_M \to \langle \varnothing \rangle_N$. Similarly, we obtain a homomorphism $h\colon \langle \varnothing \rangle_N \to \langle \varnothing \rangle_M$. Since homomorphisms out of a structure generated by $\varnothing$ are unique, $g \circ h$ and $h \circ g$ are identity maps, so these homomorphisms are inverses, and $\langle \varnothing \rangle_M \cong \langle \varnothing \rangle_N$.

(2)⇒(3): Since T is satisfiable, it has a model M. Let $A = \langle \varnothing \rangle_M$. For any $N \models T$, $A \cong \langle \varnothing \rangle_N$, so A embeds in N.

(3)⇒(1): Since T is satisfiable, it has a model M. Let φ be a sentence. Then either $M \models \varphi$ or $M \models \neg\varphi$; without loss of generality, we assume $M \models \varphi$. We claim that $T \models \varphi$.

So let $N \models T$. By our assumption, there is a structure A with embeddings $g\colon A \to M$ and $h\colon A \to N$. By quantifier elimination, φ is equivalent to a quantifier-free sentence, and by Theorem 6.7, $M \models \varphi$ implies $N \models \varphi$. □

Corollary 6.14. *If $\mathcal{L}$ has no 0-ary function symbols (constant symbols) or 0-ary relation symbols (proposition symbols), then every satisfiable theory with quantifier elimination is complete.*

Proof. If $\mathcal{L}$ has no 0-ary symbols, then there is a unique empty $\mathcal{L}$-structure: each n-ary function symbol is interpreted as the unique empty function $\varnothing^n \to \varnothing$, and each n-ary relation symbol is interpreted as the unique empty relation on $\varnothing^n$. This unique empty structure is $\langle \varnothing \rangle_M$ for every model $M \models T$. Thus, if T has quantifier elimination, then T is complete, by Proposition 6.13.

Another way of proving this is to note that if T has quantifier elimination, then every sentence is T-equivalent to a quantifier-free sentence. But if there are no 0-ary symbols, the only quantifier-free sentences are $\top$ and $\bot$. So for every sentence φ, either $T \models \varphi \leftrightarrow \top$ or $T \models \varphi \leftrightarrow \bot$, which are equivalent to $T \models \varphi$ and $T \models \neg\varphi$, respectively. □

It follows from Corollary 6.14 that the theories T_{Inf} and DLO from Examples 6.11 and 6.12 are complete, since their languages contain no constant symbols or proposition symbols.

6.3 Algebraically closed fields

Lecture 17: 11/10

We will now focus on applying our work so far to the theory ACF of algebraically closed fields, proving quantifier elimination and deriving some consequences. We work in the language of rings, $\mathcal{L}_{\text{Ring}} = \{0, 1, +, -, \cdot\}$. Let T_{Ring} be the theory of rings, and let T_{Field} be the theory of fields.

For any polynomial $p \in \mathbb{Z}[x_1, \ldots, x_n]$, there is an $\mathcal{L}_{\text{Ring}}$-term $t_p(x)$ in context $x = \{x_1, \ldots, x_n\}$ which is T_{Ring}-equivalent to p in the sense that in any ring R, for any $r \in R^x$, $t_p^R(r) = p(r)$. Conversely, it is straightforward to show by induction that every $\mathcal{L}_{\text{Ring}}$-term $t(x)$ in context $x = \{x_1, \ldots, x_n\}$ is T_{Ring}-equivalent to a polynomial in $\mathbb{Z}[x_1, \ldots, x_n]$. Similarly, if R is a ring and $A \subseteq R$, letting $S = \langle A \rangle$, every $\mathcal{L}_{\text{Ring}}(A)$-term is T_{Ring}-equivalent to a polynomial in $S[x_1, \ldots, x_n]$ and vice versa. For this reason, when working with rings, we will conflate polynomials with terms.

Every atomic $\mathcal{L}_{\text{Ring}}$-formula $\varphi(x)$ has the form $p(x) = q(x)$, where p, q are polynomials. This is T_{Ring}-equivalent to the formula $(p - q)(x) = 0$, so every atomic formula is T_{Ring}-equivalent to a polynomial equation $f(x) = 0$ (over $\mathbb{Z}$, or over $S = \langle A \rangle$ if we allow parameters).

For $n \in \mathbb{N}_{>0}$, an important example of an atomic sentence is χ_n, which asserts that the ring has characteristic dividing n:

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

We can axiomatize the class of field of characteristic p (for prime p) by $T_{\text{Field}} \cup \{\chi_p\}$, and the class of fields of characteristic 0 by $T_{\text{Field}} \cup \{\neg\chi_p \mid p \text{ prime}\}$.

The theory ACF of algebraically closed fields consists of the theory of fields T_{Field}, together with a sentence φ_d for every degree $d \geq 1$, expressing that every monic polynomial of degree d with coefficients in the field has a root in the field:

$$\forall a_0 \ldots \forall a_{d-1} \exists x \, (x^d + a_{d-1}x^{d-1} + \cdots + a_1 x + a_0 = 0).$$

We write ACF_p for $\text{ACF} \cup \{\chi_p\}$ when p is prime, and we write ACF_0 for $\text{ACF} \cup \{\neg\chi_p \mid p \text{ prime}\}$.

Theorem 6.15. ACF *has quantifier elimination.*

Proof. We use test (5) in Theorem 6.10. So suppose K_1 and K_2 are algebraically closed fields, $\varphi(x)$ is a primitive formula, and $a \in K_1^x$ such that $K_1 \models \varphi(a)$. Let $A = \langle a \rangle_{K_1}$, and let $f \colon A \to K_2$ be an embedding. We would like to show that $K_2 \models \varphi(f(a))$.

Note that since A is a subring of a field, it is an integral domain. Let F be the subfield of K_1 generated by A, i.e., $F = \{ab^{-1} \mid a \in A, b \in A \setminus \{0\}\}$. Then $F \cong \text{Frac}(A)$, the field of fractions of A, and the embedding f extends to an embedding $g \colon F \to K_2$, by $g(ab^{-1}) = f(a)f(b)^{-1}$. Let $F' = g(F)$, which is the subfield of K_2 generated by $f(A)$. For a polynomial $p \in F[x]$, write p^g for the image of p under the isomorphism $F[x] \cong F'[x]$ induced by g.

Write the primitive formula $\varphi(a)$ as $\exists y \bigwedge_{i=1}^n \varphi_i(a, y)$. Since each φ_i is atomic or negated atomic, we may assume that each φ_i is $p_i(y) = 0$ or $p_i(y) \neq 0$, where $p_i \in F[y]$. Then we can write $\varphi(a)$ as $\exists y \bigwedge_{i=1}^m p_i(y) = 0 \wedge \bigwedge_{j=1}^n q_j(y) \neq 0$.

Similarly, $\varphi(f(a))$ is $\exists y \bigwedge_{i=1}^m p_i^g(y) = 0 \wedge \bigwedge_{j=1}^n q_j^g(y) \neq 0$.

Case 1: At least one of the p_i is non-zero. Let $b \in K_1$ be a witness to the existential quantifier. Then b is algebraic over F. Let $m \in F[y]$ be the minimal polynomial of b. Then for each $1 \leq i \leq m$, $m \mid p'$, and for each $1 \leq j \leq n$, $m \nmid q_j$.

Since K_2 is algebraically closed, we can pick a root $b' \in K_2$ of m^g. Since m is irreducible, so is m^g, so m^g is the minimal polynomial of b'. Then $m^g \mid (p_i)^g$ for all $1 \leq i \leq m$ and $m^g \nmid (q_j)^g$ for all $1 \leq j \leq n$, so b' witnesses $K_2 \models \varphi(f(a))$.

Case 2: Each p_i is the zero polynomial (or $m = 0$). In this case, we only need to find $b' \in K_2$ which is not a root of any of the finitely many polynomials $(q_j)^g$. Since each polynomial has only finitely many roots in K_2 and K_2 is infinite (being algebraically closed), we can find a witness $b' \in K_2$ for the existential quantifier.

In either case, $K \models \varphi(f(a))$, as desired. □

The theory ACF is not complete, because it does not determine the characteristic: for any prime p, ACF $\not\models \chi_p$ and ACF $\not\models \neg\chi_p$. But as a first application of quantifier elimination, we will show that the theories ACF_p and ACF_0 are complete.

Note that only very basic field theory was used in the proof. Nevertheless, quantifier elimination has some very non-trivial consequences, as we will now see. We'll begin by explaining a reframing of quantifier elimination in the language of (classical) algebraic geometry.

Let K be an algebraically closed field. Fixing $m \in \omega$, let $S \subseteq K[x_1, \ldots, x_m]$ be a set of polynomials. Then we define

$$V(S) = \{(a_1, \ldots, a_m) \in K^m \mid f(a_1, \ldots, a_m) = 0 \text{ for all } f \in S\}.$$

A set of the form $V(S) \subseteq K^m$ is called a **algebraic set** or a **Zariski-closed set** (these are the closed sets of the Zariski topology on K^m).

It is easy to see that for any $S \subseteq K[x_1, \ldots, x_m]$, if $\langle S \rangle$ is the ideal generated by S, then $V(\langle S \rangle) = V(S)$. Further, the Hilbert Basis Theorem says that $K[x_1, \ldots, x_m]$ is Noetherian, i.e., every ideal is finitely generated. Thus, for every algebraic set $X \subseteq K^m$, there are finitely many polynomials $p_1, \ldots, p_n$ such that $X = V(p_1, \ldots, p_n)$.

Given algebraic sets $X \subseteq K^m$ and $Y \subseteq K^n$, a **polynomial map** $X \to Y$ is an n-tuple $f = (f_1, \ldots, f_n)$ with each $f_i \in K[x_1, \ldots, x_m]$ such that for all $a = (a_1, \ldots, a_m) \in X$, $f(a) = (f_1(a_1, \ldots, a_m), \ldots, f_n(a_1, \ldots, a_m)) \in Y$.

It is not true in general that the image of an algebraic set under a polynomial map is algebraic. For example, the image of the algebraic set $V(xy - 1) \subseteq K^2$ under the coordinate projection $K^2 \to K$ (defined by the polynomial x) is $K \setminus \{0\}$, which is not algebraic (the proper algebraic subsets of K are all finite). However, in this case we do get the complement of an algebraic set.

A subset $Y \subseteq K^n$ is **constructible** if it is a finite Boolean combination of algebraic sets.

Corollary 6.16 (Chevalley's theorem)**.** *Let $X \subseteq K^m$ be an algebraic set, and let $f\colon X \to K^n$ be a polynomial map. Then the image $f(X) \subseteq K^n$ is a constructible set.*

Proof. Let $X = V(p_1, \ldots, p_k)$. We can define $f(X)$ by the following formula:

$$\exists x \left(\bigwedge_{i=1}^{k} p_k(x) = 0 \wedge f(x) = y \right).$$

This is an $\mathcal{L}_{\text{Ring}}$-formula with parameters from K for the coefficients of the polynomials p_k and f_j. Note that f is an n-tuple of polynomials in m variables, x is an m-tuple of variables, y is an n-tuple of variables, and $f(x) = y$ is shorthand for $\bigwedge_{j=1}^{n} f_j(x) = y_j$.

By quantifier elimination, this formula is equivalent to a quantifier-free formula, $\varphi(y)$. This $\varphi(y)$ is a Boolean combination of polynomial equations $q(y) = 0$ with $q \in K[y_1, \ldots, y_n]$, so it defines a constructible subset of K^n. □

The theory ACF is not complete, since it does not determine the characteristic. But it follows from QE that we obtain a complete theory once we fix the characteristic.

Corollary 6.17. *Let p be prime or 0. Then ACF_p is complete.*

Proof. Since $\mathrm{ACF} \subseteq \mathrm{ACF}_p$, ACF_p has quantifier elimination. If $p = 0$, we take $A = \mathbb{Q}$, and if p is prime, we take $A = \mathbb{F}_p$, the finite field of order p. Then A embeds in any field of characteristic p, and in particular in every model of ACF_p, so ACF_p is complete, by Proposition 6.13. □

If a theory T has a computably enumerable axiomatization (as the theories ACF_p for p prime or 0 do), then as soon as we know T is complete, we have an algorithm for deciding which sentences are entailed by T (we say that T is a **decidable** theory). For any sentence φ, start searching for a formal proof $T \vdash \varphi$, and simultaneously start searching for a formal proof $T \vdash \neg\varphi$. Since T is complete, one of these searches eventually terminates. (Note that I have no defined formal proof in this class).

But the algorithm described above is hopelessly inefficient. In many situations, a more efficient algorithm can be found via **effective quantifier elimination**, i.e. an algorithm for finding a quantifier-free formula which is T-equivalent to a given formula. Applying this algorithm to an arbitrary sentence φ produces a quantifier-free sentence ψ which is T-equivalent to it. But a quantifier-free sentence ψ is just a Boolean combination of atomic sentences, the truth value of which can usually be easily checked.

The proof we gave above that ACF has quantifier elimination was entirely nonconstructive. An effective quantifier elimination algorithm exists for ACF;

unfortunately, it is also very inefficient in the worst case (doubly exponential running time in the size of the input).

We now dwell on some consequences of completeness for the theories ACF_p. In the following, the equivalence (1) $\Leftrightarrow$ (2) is a version of the "Lefschetz principle" that all algebraically closed fields of characteristic 0 have the same algebraic properties, while the equivalence (2) $\Leftrightarrow$ (4) makes precise an intuition that the behavior of algebraically closed fields of characteristic 0 is the "limit as p goes to ∞" of the behavior of algebraically closed fields of characteristic p.

Corollary 6.18 (Transfer principle for algebraically closed fields)**.** *Let φ be a sentence in the language of rings. The following are equivalent:*

(1) Some algebraically closed field of characteristic 0 *satisfies φ.*

(2) Every algebraically closed field of characteristic 0 *satisfies φ.*

(3) For all but finitely many primes p, some algebraically closed field of characteristic p satisfies φ.

(4) For all but finitely many primes p, every algebraically closed field of characteristic p satisfies φ.

Proof. (1)$\Rightarrow$(2): If $M \models \mathrm{ACF}_0$ and $M \models \varphi$, then $\mathrm{ACF}_0 \not\models \neg\varphi$, so $\mathrm{ACF}_0 \models \varphi$ by completeness.

(2)$\Rightarrow$(4): We have $\mathrm{ACF}_0 \models \varphi$, so by compactness there are finitely many primes $p_1, \ldots, p_n$ such that $\mathrm{ACF} \cup \{p_i \neq 0 \mid 1 \leq i \leq n\} \models \varphi$. For any prime q not among these finitely many exceptions, any algebraically closed field of characteristic q satisfies $\mathrm{ACF} \cup \{p_i \neq 0 \mid 1 \leq i \leq n\}$, hence satisfies φ.

(4)$\Rightarrow$(3): Trivial.

(3)$\Rightarrow$(1): Assume for contradiction that (1) fails. Then every algebraically closed field of characteristic 0 satisfies $\neg\varphi$. By (2)$\Rightarrow$(4) for $\neg\varphi$, we have that for all but finitely many primes p, every algebraically closed field of characteristic p satisfies $\neg\varphi$. This contradicts (3) (since there are infinitely many primes!). $\square$

Here is a nice application of the transfer principle.

Theorem 6.19 (Ax–Grothendieck)**.** *Any injective polynomial map $\mathbb{C}^n \to \mathbb{C}^n$ is surjective.*

The statement is obviously true if we replace $\mathbb{C}$ by a finite fields (since any injective function from a finite set to itself is surjective). We will use model theory to transfer the result to $\mathbb{C}$. The key observation is the following general statement.

Theorem 6.20. *Let φ be a $\forall\exists$-sentence in $\mathcal{L}_{Ring}$ which is true in every finite field. Then φ is true in every algebraically closed field.*

Proof. Fix a prime number p, and let $K = \overline{\mathbb{F}_p}$, the algebraic closure of the prime field $\mathbb{F}_p$. Then $K \cong \varinjlim \mathbb{F}_{p^n}$, which is a direct limit in which the index set $\{n \in \omega \mid n > 0\}$ is ordered by the divisibility order.

We assumed that φ is true in every finite field, and since it is inductive ($\forall\exists$), it is preserved in direct limits, so $K \models \varphi$, and by completeness, $\mathrm{ACF}_p \models \varphi$.

We have shown that φ holds in all algebraically closed fields of finite characteristic. It follows that it holds in all algebraically closed fields of characteristic 0 by Corollary 6.18. □

Proof of Theorem 6.19. For each n and each degree d, we can express the statement of the theorem restricted to polynomial maps $\mathbb{C}^n \to \mathbb{C}^n$ such that each polynomial has degree at most d, by the following sentence φ_d:

$$\forall f (\forall x\, \forall x' ((f(x) = f(x')) \to (x = x')) \to \forall y\, \exists x\, (f(x) = y))$$

There are a few things to explain here.

(1) The quantifier $\forall f$ means to quantify over all polynomial maps $\mathbb{C}^n \to \mathbb{C}^n$ of degree at most d. Such a polynomial map is an n-tuple of polynomials in $\mathbb{C}[x_1, \dots, x_n]$, each of degree at most d. To do this, we can quantify over the coefficients of these polynomials. So really we have $n(d+1)$ quantifiers, over variables $(a_i^j)_{1\leq i\leq n, 0\leq j\leq d}$, which specify polynomials $p_i = a_i^0 + a_i^1 x + \cdots + a_i^d x^d$ for $1 \leq i \leq n$.

(2) Similarly, the quantifiers over x, x' and y are over n-tuples, e.g., $x = (x_i)_{1\leq i\leq n}$.

(3) The formulas $f(x) = f(x')$, $x = x'$, and $f(x) = y$ are really conjunctions of n atomic formulas, e.g. $f(x) = f(x')$ is $\bigwedge_{i=1}^n (p_i(x_1, \dots, x_n) = p_i(x_1', \dots, x_n'))$.

(4) φ_d is logically equivalent to a $\forall\exists$-sentence. We have to be slightly careful about the presence of $\to$, because $\varphi \to \psi$ is shorthand for $\neg\varphi \vee \psi$, and $\neg$ is not allowed in the construction of $\forall\exists$-formulas. But we can rewrite φ_d as follows:

$$\forall f\, \forall y\, (\exists x\, \exists x'\, (f(x) = f(x') \wedge (x \neq x'))) \vee \exists x\, (f(x) = y)).$$

Now for all d and any finite field F, we have $F \models \varphi_d$, since an injective map from a finite set to itself is always surjective. By Theorem 6.20, $\mathbb{C} \models \varphi_d$ for all d, as desired. □

It is interesting to note that while it is also true that every surjective function from a finite set to itself is injective, the converse of the Ax–Grothendieck theorem fails. For example, $x \mapsto x^2$ is a surjective polynomial map $\mathbb{C} \to \mathbb{C}$ which is not injective. The proof fails in this case because the sentences asserting "every surjective polynomial map is injective" are not equivalent to $\forall\exists$-sentences. They are $\forall\exists\forall$.

Lecture 18: 11/12

We now return to some application that rely on the full strength of quantifier elimination, not just completeness.

Definition 6.21. A theory T is **strongly minimal** if for every model $M \models T$ and every $\mathcal{L}(M)$-formula $\varphi(x)$ with a single free variable x, $\varphi(M)$ is either finite or cofinite.

Theorem 6.22. ACF *is strongly minimal.*

Proof. Let $K \models \mathrm{ACF}$, and let $\varphi(x)$ be an $\mathcal{L}(K)$-formula with one free variable. By QE, $\varphi(x)$ is equivalent to a quantifier-free $\mathcal{L}(K)$-formula $\psi(x)$.

Every atomic $\mathcal{L}(K)$-formula $\chi(x)$ is equivalent to $p(x) = 0$, where $p \in K[x]$, and hence $\chi(K)$ is finite (if p is non-zero) or K (if p is 0). In particular, $\chi(K)$ is finite or cofinite. Since the set of finite and cofinite subsets of K form a Boolean algebra (they are closed under intersection, union, and complement and contain $\varnothing$ and K), and $\psi(x)$ is a Boolean combination of atomic $\mathcal{L}(K)$-formulas, $\varphi(K) = \psi(K)$ is finite or cofinite. $\square$

Strong minimality is a condition on formulas in one free variable (definable subsets of M), but it has powerful consequences for all formulas (definable subsets of M^n for all n) and on the class of models of T. I'll just give an indication of this.

Theorem 6.23. *Let T be a strongly minimal theory. Let $\varphi(x, y)$ be a formula where y is an arbitrary finite variable context and x is a single variable. Then there exists $N \in \omega$ such that for all $M \models T$ and all $b \in M^y$, either $|\varphi(M, b)| < N$ or $|M \setminus \varphi(M, b)| < N$.*

Proof. Consider the following partial type $\Sigma(y)$:

$$\{\exists^{\geq n} x\, \varphi(x, y) \wedge \exists^{\geq n} x\, \neg\varphi(x, y) \mid n \in \omega\}.$$

If $M \models T$ and $b \in M^y$ realizes $\Sigma(y)$, then $\varphi(M, b)$ is infinite and coinfinite, contradicting strong minimality of T.

Thus $\Sigma(y)$ is not satisfiable relative to T, and by compactness there is some $N \in \omega$ such that $\exists^{\geq N} x\, \varphi(x, y) \wedge \exists^{\geq N} x\, \neg\varphi(x, y)$ is unsatisfiable in a model of T. It follows that for all $M \models T$ and all $b \in M^y$, either $|\varphi(M, b)| < N$ or $|M \setminus \varphi(M, b)| < N$. $\square$

For example, consider the case when $\varphi(x, y)$ is $\bigwedge_{i=1}^{n} p_i(x, y) = 0$, a system of polynomial equations defining the algebraic set $V(p_1, \ldots, p_n)$. The theorem implies that if $\pi\colon K^{n+1} \to K^n$ is the projection map which forgets the x-coordinate, then there is a uniform upper bound on the size of finite fibers – and in fact, this bound is uniform across all algebraically closed fields, depending only on the syntactic form of φ, i.e., the number of polynomials and their degrees.

Here's a theorem that I won't prove, for time reasons.

Theorem 6.24. *Let T be a strongly minimal complete theory. If $\kappa > |\mathcal{L}|$, then T is κ-**categorical**: Up to isomorphism, T has a unique model of cardinality κ.*

It follows that, since each theory ACF_p is complete and strongly minimal in a countable language, any two uncountable algebraically closed fields of the same characteristic and cardinality are isomorphic.

Finally, we return to model completeness and existentially closed models. Since ACF eliminates quantifiers, it follows immediately that ACF is model complete, and hence by Theorem 6.3 that every algebraically closed field is existentially closed (in the class of models of ACF). All the way back in Example 4.9, we proved that a ring which is PEC in the class of non-zero rings is an algebraically closed field, but we left open whether the converse is true. We can now prove that it is:

Theorem 6.25. *Let R be a non-zero ring. The following are equivalent:*

(1) R is an algebraically closed field.

(2) R is EC in the class of models of ACF*.*

(3) R is PEC in the class of non-zero rings.

Proof. (3)⇒(1): by Example 4.9.

(1)⇒(2): By Theorem 6.3, since ACF is model complete.

(2)⇒(3): Suppose K is EC in the class of models of ACF (so in particular, K is an algebraically closed field). Let $h\colon K \to R$ be a homomorphism, where R is a non-zero ring. Let $\varphi(x)$ be an $\exists^+$-formula, $a \in K^x$, and assume $R \models \varphi(h(a))$. We must show $K \models \varphi(a)$.

By Theorem 4.10, let $h'\colon R \to F$ be a homomorphism from R to a PEC non-zero ring F, and note that $F \models$ ACF by (3)⇒(1). (Alternatively, pick a maximal ideal M in R, map R to the field R/M, and embed this field in its algebraic closure.) The homomorphism h' preserves $\exists^+$-formulas, so $F \models \varphi(h'(h(a)))$. Now $h' \circ h\colon K \to F$ is a homomorphism between fields, so it is injective, and hence an embedding. Since K is EC in the class of models of ACF, $h' \circ h$ reflects φ, so $K \models \varphi(a)$, as desired. □

Note that the definition of algebraically closed field only requires that polynomials in one variable have roots, while the PEC condition refers to arbitrary $\exists^+$-formulas, i.e., roots of systems of polynomials in multiple variables. The content of (1) ⇒ (3) in the above theorem is essentially Hilbert's Nullstellensatz. I state it here in its "weak" form, because this is the one that follows most immediately. It is easy (with a bit of commutative algebra) to derive the other forms of the Nullstellensatz from this one, or to give model-theoretic proofs using again the fact that models of ACF are EC.

Theorem 6.26 ((Weak) Nullstellensatz)**.** *Let K be an algebraically closed field. For all $p_1, \dots, p_n \in K[x_1, \dots, x_m]$, if $1 \notin (p_1, \dots, p_n)$ (i.e., the p_i do not generate the unit ideal in the polynomial ring), then $V(p_1, \dots, p_n)$ is non-empty.*

Proof. Let $R = K[x_1, \dots, x_m]$. Let $h\colon K \to R$ be homomorphism which includes K in R as the constant polynomials. Let $I = (p_1, \dots, p_n)$, and let $q\colon R \to R/I$ be the quotient homomorphism. Since I is not the unit ideal, R/I is a non-zero ring. In R/I, $(q(x_1), \dots, q(x_n))$ satisfies the formula $\bigwedge_{i=1}^n p_i(x) = 0$. Indeed, for all $1 \leq i \leq n$, $p_i(q(x_1), \dots, q(x_n)) = q(p_i) = 0$, since $p_i \in I$. So $R/I \models \exists x \bigwedge_{i=1}^n p_i(x)$. Since K is PEC in the class of non-zero rings, $q \circ h$ reflects this formula, and $K \models \exists x \bigwedge_{i=1}^n p_i(x)$. That is, $V(p_1, \dots, p_n) \subseteq K^m$ is non-empty. □

It is a theorem of Macintyre that for an infinite field K, if $\mathrm{Th}(K)$ has QE (in $\mathcal{L}_{\mathrm{Ring}}$), then K is algebraically closed. However, there are many non-algebraically closed fields K such that $\mathrm{Th}(K)$ has quantifier elimination in a mild expansion of the language, or such that $\mathrm{Th}(K)$ is model complete (in $\mathcal{L}_{\mathrm{Ring}}$).

One example is the field of p-adic numbers $\mathbb{Q}_p$, which is model complete in $\mathcal{L}_{\mathrm{Ring}}$ and has quantifier elimination in the language $\mathcal{L}_{\mathrm{Ring}} \cup \{^{-1}, V, (P_n)_{n\in\omega}\}$, where $^{-1}$ is a unary function symbol for multiplicative inverse (with $0^{-1} = 0$), V is a unary relation symbol picking out the valuation ring $\mathbb{Z}_p$, and each P_n is a unary predicate picking out the set of n^{th} powers. This is also due to Macintyre.

Another example is the real field, which Tarski showed has quantifier elimination if we include the order in the language.

Example 6.27. Let $\mathrm{RCF} = \mathrm{Th}(\mathbb{R})$ in $\mathcal{L}_{\mathrm{Ring}}$ (RCF stands for the theory of real closed fields). Let $\mathrm{RCF}_< = \mathrm{Th}(\mathbb{R})$ in $\mathcal{L}_{\mathrm{Ring},<} = \mathcal{L}_{\mathrm{Ring}} \cup \{<\}$. It is a theorem of Tarski that $\mathrm{RCF}_<$ has quantifier elimination and is complete. I claim that it follows that RCF is model complete.

Note that the atomic formula $x < y$ is equivalent in $\mathbb{R}$ to $(x \neq y) \wedge \exists z\,(z^2 = y - x)$, and its negation $\neg(x < y)$ is equivalent to $y \geq x$, which is equivalent to $\exists z\,(z^2 = x - y)$.

Now let $\varphi(x)$ be an arbitrary $\mathcal{L}_{\mathrm{Ring}}$-formula. By QE for $\mathrm{RCF}_<$, $\varphi(x)$ is equivalent to a quantifier-free $\mathcal{L}_{\mathrm{Ring},<}$-formula $\psi(x)$, so $\varphi(R) = \psi(R)$. Write $\psi(x)$ in disjunctive normal form as a disjunction of conjunctions of literals. Now each literal involving the symbol $<$ can be replaced by an $\exists$-formula in $\mathcal{L}_{\mathrm{Ring}}$, as above, resulting in an $\exists$-formula $\theta(x)$ in $\mathcal{L}_{\mathrm{Ring}}$. Note it was important here that *both* $x < y$ and $\neg(x < y)$ were equivalent to $\exists$-formulas. Now $\varphi(\mathbb{R}) = \theta(\mathbb{R})$, so by completeness of RCF, $\mathrm{RCF} \models \forall x\,\varphi(x) \leftrightarrow \theta(x)$, i.e., these formulas are RCF-equivalent. We have proved that every $\mathcal{L}_{\mathrm{Ring}}$-formula is RCF-equivalent to an $\exists$-formula, so RCF is model complete.

For example, the formula $\varphi(a, b, c)\colon \exists x\, ax^2 + bx + c$ is $\mathrm{RCF}_<$-equivalent to $b^2 - 4ac \geq 0$, which is equivalent to $\exists z\, z^2 = b^2 - 4ac$.

We can now mimic the proof of the Nullstellensatz to obtain a Nullstellensatz for real closed fields. There is an obstruction here: not every non-zero ring admits a homomorphism to a model of RCF, so we need to understand conditions on an ideal I such that R/I admits such a homomorphism. Call an ideal I **real** if whenever $\sum_{i=1}^{n} p_i^2 \in I$, each $p_i \in I$. Every ideal I is contained in a smallest real ideal, called its **real radical**, and R/I admits a homomorphism to a model of RCF if and only if I is a proper real ideal.

We then obtain the following statement: For all $p_1, \ldots, p_n \in \mathbb{R}[x_1, \ldots, x_m]$, if the real radical of $(p_1, \ldots, p_n)$ is proper, then $V(p_1, \ldots, p_n) \subseteq \mathbb{R}^m$ is non-empty. The condition that the real radical is proper is equivalent to: for all $q \in (p_1, \ldots, p_n)$ and all $h_1, \ldots, h_k \in \mathbb{R}[x_1, \ldots, x_m]$,

$$q + \sum_{i=1}^{k} h_i^2 \neq -1.$$

6.4 Model companions

Lecture 19: 11/17

As we have seen, some theories of mathematical interest happen to be model complete in a natural language $\mathcal{L}$ and we can often prove this by proving a quantifier elimination result (in $\mathcal{L}$ or a mild expansion). Another way of obtaining a model complete theory is to start with a theory (usually a universal or inductive theory which is far from being complete) and pass to a related model complete theory called the model companion.

Given a theory T, we write $T_\forall$ for the set of universal consequence of T, i.e., the set of all $\forall$-sentences φ such that $T \models \varphi$.

Lemma 6.28. *Let M be an $\mathcal{L}$-structure. Then $M \models T_\forall$ if and only if there exists $N \models T$ and an embedding $f\colon M \to N$.*

Proof. Suppose $M \models T_\forall$. Let $T' = T \cup \mathrm{Diag}(M)$. If T' is unsatisfiable, by compactness there exists $\varphi(m)$, a finite conjunction of literals in $\mathrm{Diag}(M)$, such that $T \cup \{\varphi(m)\}$ is unsatisfiable. But then $T \models \forall x\, \neg\varphi(x)$, so $\forall x\, \neg\varphi(x) \in T_\forall$, contradicting $M \models T_\forall$ and $M \models \varphi(m)$. Letting $N \models T'$, we have $N \models T$ and an embedding $f\colon M \to N$.

Conversely, if $f\colon M \to N$ is an embedding and $N \models T$, then $N \models T_\forall$, and since f reflects $\forall$-sentences, $M \models T_\forall$. □

Example 6.29. The class of $\mathcal{L}_{\mathrm{Ring}}$-structures which embed in an algebraically closed field is the class of integral domains. Thus $\mathrm{ACF}_\forall = (T_{\mathrm{Field}})_\forall = T_{\mathrm{ID}}$, the theory of integral domains.

Every linear order L embeds in a dense linear order without endpoints. For example, we can take $x \mapsto (x, 0) \in L \times \mathbb{Q}$ ordered lexicographically. Similarly, every linear order L embeds in a discrete linear order without endpoints. For example, we can take $x \mapsto (x, 0) \in L \times \mathbb{Z}$ ordered lexicographically. If follows that if T is the theory of discrete linear orders without endpoints, then $T_\forall = \mathrm{DLO}_\forall = T_{\mathrm{LO}}$, the theory of linear orders.

Definition 6.30. Theories T and T' are **companions** if every model of T embeds in a model of T' and vice versa.

Lemma 6.31. *T and T' are companions if and only if $T_\forall \equiv T'_\forall$.*

Proof. Suppose T and T' are companions. Let $M \models T_\forall$. Then by Lemma 6.28, M embeds in $N \models T$, and since T and T' are companions, N embeds in $N' \models T'$. Thus M embeds in a model of T', so $M \models T'_\forall$. Thus $T_\forall \models T'_\forall$. The same argument with T and T' reversed shows $T_\forall \equiv T'_\forall$.

Conversely, suppose $T_\forall \equiv T'_\forall$. Let $M \models T$. Then $M \models T_\forall$, so $M \models T'_\forall$, so M embeds in a model of T' by Lemma 6.28. The same argument with T and T' reversed shows T and T' are companions. □

It is straightforward to check from the definition, and obvious from Lemma 6.31, that companionship is an equivalence relation on theories.

Definition 6.32. Let T be a theory. A theory T^* is a **model companion** of T if T and T^* are companions and T^* is model complete.

Example 6.33. ACF is a model companion of the theory of fields and the model companion of the theory of integral domains.

DLO is a model companion of the theory of discrete linear orders without endpoints.

Lemma 6.34. *Suppose T and T' are companions, T is an inductive theory, and T' is a model complete theory. Then $T' \models T$.*

Proof. Let $M_0 \models T'$. Given $M_i \models T'$, we can find an embedding $f_i \colon M_i \to N_i \models T$ and an embedding $g_i \colon N_i \to M_{i+1} \models T'$, by companionship. Thus we build a diagram:

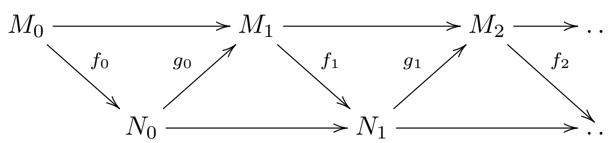

Since T is inductive, $\varinjlim N_i \models T$. But $\varinjlim M_i \cong \varinjlim N_i$, so $\varinjlim M_i \models T$.

Since T' is model complete, T' is also inductive by Theorem 6.4. Thus $\varinjlim M_i \models T'$. Now M_0 embeds in $\varinjlim M_i$, and because T' is model complete, this embedding is elementary. So since $\varinjlim M_i \models T$, also $M_0 \models T$. $\square$

It follows from Lemma 6.34 that if T has a model companion T^*, then T^* is the maximal inductive theory in the companionship class of T, in the sense that it entails every other inductive theory in the class. Dually, $T_\forall$ is the minimal theory in the companionship class of T, since every theory in the class entails $T_\forall$.

Theorem 6.35. *Any theory has at most one model companion up to equivalence.*

Proof. Suppose T^* and T^{**} are both model companions of T. In particular, T^* and T^{**} are companions. Since T^* and T^{**} are both model complete, they are both inductive by Theorem 6.4. By Lemma 6.34, $T^* \models T^{**}$ and $T^{**} \models T^*$, so these theories are equivalent. $\square$

In the case of an inductive theory T, if T has a model companion T^*, then by Lemma 6.34, every model of T^* is a model of T. Our goal now is to show that the models of T^* are exactly the existentially closed models of T. For any theory T, define

$$\mathrm{EC}(T) = \{M \models T \mid M \text{ is EC in the class of models of } T\}.$$

The following Lemma has essentially the same content as Lemma 4.13, which was a key step in our proof of the compactness theorem. Now that we have compactness, the proof is much easier.

Lemma 6.36. *Let T be an inductive theory. Then* $\mathrm{EC}(T) = \mathrm{EC}(T_\forall)$.

Proof. Assume first that $M \in \mathrm{EC}(T)$. Since $M \models T$, also $M \models T_\forall$. Now let $f\colon M \to N$ be an embedding with $N \models T_\forall$. Since T and $T_\forall$ are companions, there is an embedding $g\colon N \to M' \models T$. Let $\varphi(x)$ be an $\exists$-formula and $a \in M^x$. If $N \models \varphi(f(a))$, then $M' \models \varphi(g(f(a)))$, and since M is EC in models of T, $M \models \varphi(a)$. Thus $M \in \mathrm{EC}(T_\forall)$.

Conversely, assume $M \in \mathrm{EC}(T_\forall)$. Write T as a set of $\exists$-sequents. Let $\varphi \vdash_x \psi \in T$. Assume that $M \models \varphi(a)$ for some $a \in M^x$. We show $M \models \psi(a)$.

Since T and $T_\forall$ are companions, let $f\colon M \to N$ be an embedding with $N \models T$. Then also $N \models T_\forall$, so f preserves and reflects $\exists$-formulas, since M is EC. Thus $N \models \varphi(f(a))$, and since $N \models T$, $N \models \psi(f(a))$, so $M \models \psi(a)$.

Since M is EC in the class of models of $T_\forall$ and $M \models T$, M is clearly EC in the smaller class of models of T, so $M \in \mathrm{EC}(T)$. □

Corollary 6.37. *Suppose T and T' are inductive theories which are companions. Then $\mathrm{EC}(T) = \mathrm{EC}(T')$.*

Proof. By Lemma 6.36, $\mathrm{EC}(T) = \mathrm{EC}(T_\forall) = \mathrm{EC}(T')$. □

Theorem 6.38. *Let T be an inductive theory. Then T has a model companion if and only if the class $\mathrm{EC}(T)$ of EC models of T is elementary. In this case, the model companion T^* is $\mathrm{Th}(\mathrm{EC}(T))$.*

Proof. Suppose T has a model companion T^*. I claim that T^* axiomatizes $\mathrm{EC}(T)$, so $\mathrm{EC}(T)$ is elementary.

Suppose $M \models T^*$. Since T^* is model complete, it is inductive. Since T and T^* are companions, $\mathrm{EC}(T) = \mathrm{EC}(T^*)$ by Corollary 6.37. But since T^* is model complete, every model of T^* is in $\mathrm{EC}(T^*)$, so the class of models of T^* is exactly $\mathrm{EC}(T)$.

Now suppose $\mathrm{EC}(T)$ is elementary, axiomatized by T'. We show that T' is a model companion of T. First, we show that T and T' are companions. Every model of T' is an EC model of T, in particular a model of T, and it embeds in itself. So let $M \models T$. Since T is inductive, by Theorem 4.10, M embeds in an EC model of T, which is a model of T'.

Finally, we show that T' is model complete. If $M \models T'$, then M is EC in the class of models of T. Then M is also EC in the smaller class of models of T'. Since every model of T' is EC in the class of models of T', T' is model complete by Theorem 6.3. □

Example 6.39. T_{Field} is an inductive theory, and its model companion $\mathrm{ACF} = \mathrm{Th}(\mathrm{EC}(T_{\mathrm{Field}}))$.

Letting T be the theory of discrete linear orders without endpoints, T is not inductive, and indeed a model of DLO is not a model of T. But $T_\forall = T_{\mathrm{LO}}$ is inductive, and $\mathrm{DLO} = \mathrm{Th}(\mathrm{EC}(T_{\mathrm{LO}}))$.

Example 6.40. Let $\mathcal{L}_{\mathrm{Graph}} = \{E\}$, and let T be the theory of acyclic graphs (i.e., forests), axiomatized by:

- $\forall x\, \neg xEx$.

- $\forall x\,(xEy \to yEx)$.
- For all $n \geq 3$, $\forall x_1 \dots x_n \neg \Big(x_n E x_1 \wedge \bigwedge_{i=1}^{n-1} x_i = x_{i+1} \Big)$.

I claim that T has no model companion.

Suppose for contradiction that T has a model companion T^*. Note that T is a $\forall$-theory, hence inductive, by Theorem 6.38, T^* is the theory of existentially closed acyclic graphs.

Let $\varphi_n(x, y)$ be the formula asserting that there is no path of length $\leq n$ from x to y, and let $\Sigma(x, y) = \{\varphi_n(x, y) \mid n \in \omega\}$. Each $\varphi_n(x, y)$ is realized by the endpoints a_0 and a_{n+1} of the path graph P_n of length $(n + 1)$. Now P_n embeds in an EC $M_n \models T^*$, and since M_n is acyclic, no shorted path from a_0 to a_{n+1} can be introduced, so $M_n \models \varphi_n(a_0, a_{n+1})$. By compactness, there is a model $M \models T^*$ and $a, b \in M$ realizing Σ, i.e., there is no path from from a to b. So M is not connected.

But now we can embed M in $M' = M \cup \{c\}$ with a new vertex c connected only to a and b. Since there is no path from a to b in M, M' is acyclic, so $M' \models T$. But the inclusion $M \to M'$ fails to reflect the formula $\exists z\,(xEz \wedge zEy)$, contradicting the fact that M' is EC.

Another way to view the above example is that every EC acyclic graph is connected, but this property is incompatible with the compactness theorem. This is a typical strategy for showing that a theory has no model companion.

Lecture 20: 11/19

Example 6.41. We now show that T_{Group} has no model companion.

It is a fact that for any group G and $a, b \in G$ of the same order (finite or infintie), there exists a group extension $G \leq H$ such that a and b are conjugate by an inner automorphism in H, i.e., there exists $h \in H$ such that $hah^{-1} = b$. H is called an HNN extension of G (after Graham Higman, Bernhard Neumann, and Hanna Neumann). The construction of H is easy: presenting G by generators and relations as $G = \langle X \mid R \rangle$, let $H = \langle X, h \mid R, hah^1b^{-1} \rangle$. The non-trivial part is showing that the homomorphism $G \to H$ induced by these presentations is an embedding.

It follows that if G is an EC group and $a, b \in G$ have the same order, then a and b are conjugate by an inner automorphism in G (since the embedding of G in the HNN extension reflects the formula $\exists z\,(zxz^{-1} = y)$).

Now suppose for contradiction that the class of EC groups is axiomatizable by T^*_{Group}. We construct a model of T^*_{Group} with two elements of infinite order which are not conjugate by an inner automorphism.

Let

$$\Sigma(x, y) = \{\neg \exists z\,(zxz^{-1} = y)\} \cup \{x^n \neq e, y^n \neq e \mid n \geq 1\}.$$

By compactness, it suffices to find a model of T^*_{Group} containing two elements of order greater than n which are not conjugate by an inner automorphism. Let p and q be primes greater than n, and let G be a group with elements a and b of order p and q, respectively. Embed G in an EC group $G^* \models T^*_{\text{Group}}$. Since the embedding $f\colon G \to G'$ is injective, a and b still have order p and q in G^*. Since they have different orders, they cannot be conjugate by an automorphism.

For any $\exists$-formula $\varphi(x)$, define

$$\mathrm{Cont}^T_\varphi = \{\psi(x,y) \text{ quantifier-free} \mid T \models \forall x \forall y\, \neg(\psi(x,y) \wedge \varphi(x))\}.$$

In other words, Cont^T_φ consists of quantifier-free formulas, possibly in larger variable contexts, which contradict φ in models of T. Another way to put it is that $\varphi(x,y) \in \mathrm{Cont}^T_\varphi$ if and only if $T \models (\varphi(x) \vdash_{xy} \neg\psi(x,y))$ if and only if $T \models (\varphi(x) \vdash_x \forall y\, \neg\psi(x,y))$. Hodges calls the set of all $\forall$-formulas entailed by $\varphi(x)$ the "resultant" of φ (in analogy with the polynomial resultant in algebra). The content of the Hodges's resultant is essentially the same as my Cont^T_φ.

Lemma 6.42. *Let T be a theory. Then $M \in \mathrm{EC}(T)$ if and only if $M \models T$ and for every $\exists$-formula $\varphi(x)$ and every $a \in M^x$, either $M \models \varphi(a)$ or there exists $\psi(x,y) \in \mathrm{Cont}^T_\varphi$ and $b \in M^y$ such that $M \models \psi(a,b)$.*

Proof. First, assume $M \models T$ is EC. Let $\varphi(x)$ be an $\exists$-formula and let $a \in M^x$, and assume $M \not\models \varphi(a)$. Since M is EC, $T \cup \mathrm{Diag}(M) \cup \{\varphi(a)\}$ is unsatisfiable. By compactness there is a formula $\psi(a,b)$, a finite conjunction of literals in $\mathrm{Diag}(M)$, such that $T \cup \{\psi(a,b), \varphi(a)\}$ is unsatisfiable. Then $T \models \forall x\, \neg(\varphi(x) \wedge \psi(x,y))$, so $\psi(x,y) \in \mathrm{Cont}^T_\varphi(x)$, and $M \models \psi(a,b)$.

Conversely, let $f\colon M \to N$ be an embedding, with $M, N \models T$, let $\varphi(x)$ be an $\exists$-formula, and let $a \in M^x$. It suffices to show that if $M \not\models \varphi(a)$, then $N \not\models \varphi(f(a))$. By our hypothesis on M, there exists $\psi(x,y) \in \mathrm{Cont}^T_\varphi$ and $b \in M^y$ such that $M \models \psi(a,b)$. Then $N \models \psi(f(a), f(b))$. Since $\psi(x,y) \in \mathrm{Cont}^T_\varphi$ and $N \models T$, $N \not\models \varphi(f(a))$, as desired. □

Theorem 6.43. *Let T be an inductive theory with a model companion T^*. Then for every $\exists$-formula $\varphi(x)$, there is an $\exists$-formula $\widetilde{\varphi}(x)$ such that $T^* \models \forall x\, (\neg\varphi(x) \leftrightarrow \widetilde{\varphi}(x))$. With this notation, T^* can be axiomatized as*

$$T \cup \{\forall x\, (\varphi(x) \vee \widetilde{\varphi}(x)) \mid \varphi(x) \text{ an } \exists\text{-formula } \}.$$

Proof. Since T^* is model complete, by Theorem 6.3, for every $\exists$-formula $\varphi(x)$, $\neg\varphi(x)$ is T-equivalent to an $\exists$-formula $\widetilde{\varphi}(x)$. Note that if we write $\widetilde{\varphi}(x)$ in normal form as $\bigvee_{i=1}^k \exists y^i\, \psi_i(x, y^i)$, where each ψ_i is quantifier-free, then $T^* \models \forall x \forall y^i\, (\psi_i(x,y^i) \to \widetilde{\varphi}(x))$, so $T^* \models \forall x \forall y\, \neg(\psi_i(x,y^i) \wedge \varphi(x))$. This is a $\forall$-sentence, so since T and T^* are companions, also $T \models \forall x \forall y\, \neg(\psi_i(x,y^i) \wedge \varphi(x))$. Thus each $\psi_i(x,y^i)$ is in Cont^T_φ.

Since T is inductive, $T^* \models T$, so every model of T^* satisfies

$$T \cup \{\forall x\, (\varphi(x) \vee \widetilde{\varphi}(x)) \mid \varphi(x) \text{ an } \exists\text{-formula } \}.$$

Conversely, suppose M is a model of this theory. It suffices to show that $M \in \mathrm{EC}(T)$. Let $\varphi(x)$ be an $\exists$-formula and $a \in M^x$, and assume that $M \not\models \varphi(a)$. Then $M \models \widetilde{\varphi}(a)$, so there is some i and some $b \in M^{y^i}$ such that $M \models \psi_i(a,b)$. Since $\psi_i(x,y^i) \in \mathrm{Cont}^T_\varphi$, by Lemma 6.42, $M \in \mathrm{EC}(T)$. □

Properties of the model companion T^* can often be determined from properties of T. Here we consider "how complete" the model companion is. Note that since T^* is model complete, for any $M \models T^*$, $T^* \cup \mathrm{Diag}(M)$ is complete. And T^* has QE if and only if for any substructure $A \subseteq M \models T^*$, $T^* \cup \mathrm{Diag}(A)$ is complete. The next definition generalizes this.

Definition 6.44. Suppose T^* is a model companion of T. We say that T^* is a **model completion** of T if for every $M \models T$, $T^* \cup \mathrm{Diag}(M)$ is a complete $\mathcal{L}(M)$-theory.

The name "model completion" is rather unfortunate, since a model completion need not be complete. For example, ACF is a model completion of T_{Field}.

Definition 6.45. Let T be a theory.

1. T has the **joint embedding property (JEP)** if for all $B_1, B_2 \models T$, there exists $C \models T$ and embeddings $g_1 \colon B_1 \to C$ and $g_2 \colon B_2 \to C$.
2. $A \models T$ is an **amalgamation base** if for all $B_1, B_2 \models T$ and embeddings $f_1 \colon A \to B_1$ and $f_2 \colon A \to B_2$, there exists $C \models T$ and embeddings $g_1 \colon B_1 \to C$ and $g_2 \colon B_2 \to C$ such that $g_1 \circ f_1 = g_2 \circ f_2$.
3. T has the **amalgamation property (AP)** if every $A \models T$ is an amalgamation base.

Lemma 6.46. *Suppose T is a complete theory. Then T has JEP for elementary embeddings: for all $B_1, B_2 \models T$, there exists $C \models T$ and elementary embeddings $g_1 \colon B_1 \to C$ and $g_2 \colon B_2 \to C$.*

Proof. Let $M_1, M_2 \models T$. The empty map $M \dashrightarrow N$ is a partial elementary map, because $M_1 \equiv M_2$. By Theorem 5.19, there exists an elementary extension $M_2 \preceq N$ and an elementary embedding $M_1 \to N$. □

Lemma 6.47. *Suppose T and T' are companions. Then T has JEP if and only if T' has JEP.*

Proof. Since companionship is symmetric, it suffices to assume T' has JEP and prove T has JEP.

Let $B_1, B_2 \models T$. Since T and T' are companions, there exist embeddings $h_1 \colon B_1 \to B_1'$ and $h_2 \colon B_2 \to B_2'$ with $B_1', B_2' \models T'$. Now since T' has JEP, there are embeddings $g_1 \colon B_1' \to C'$ and $g_2 \colon B_2' \to C'$ with $C' \models T'$. Finally, by companionship again, there is an embedding $f \colon C' \to C$ with $C \models T$. So $f \circ g_1 \circ h_1$ and $f \circ g_2 \circ h_2$ witness that T has JEP. □

Theorem 6.48. *Suppose T^* is a model companion of T. Then T^* is complete if and only if T has JEP.*

Proof. Suppose T^* is complete. By Lemma 6.46, T^* has JEP. Since T and T^* are companions, by Lemma 6.47, T has JEP.

Conversely, suppose T has JEP. To show T^* is complete, it suffices to show that for any $M_1, M_2 \models T^*$, $M_1 \equiv M_2$. By Lemma 6.47, T^* has JEP. Thus there are embeddings $g_1 \colon M_1 \to N$ and $g_2 \colon M_2 \to N$ with $N \models T^*$. Since T^* is model complete, g_1 and g_2 are elementary embeddings, so $M_1 \equiv N \equiv M_2$. □

Theorem 6.49. *Suppose T^* is a model companion of T and $A \models T$. Then A is an amalgamation base for T if and only if $T^* \cup \mathrm{Diag}(A)$ is a complete $\mathcal{L}(A)$-theory.*

Proof. Write $T(A)$ for $T \cup \mathrm{Diag}(A)$ and $T^*(A)$ for $T^* \cup \mathrm{Diag}(A)$. We need to observe the following things:

(1) A model of $T(A)$ is a model B of T together with an embedding $f \colon A \to B$. If B and B' are models of $T(A)$, with embeddings $f \colon A \to B$ and $g \colon A \to B'$, an $\mathcal{L}(A)$-embedding $h' \colon B \to B'$ is the same as an $\mathcal{L}$-embedding $h \colon B \to B'$ such that $h \circ f = g$.

(2) A is an amalgamation base for T if and only if $T(A)$ has JEP. Indeed, given $f_1 \colon A \to B_1$ and $f_2 \colon A \to B_2$, we can view B_1 and B_2 as models of $T(A)$. Then by JEP, there are $\mathcal{L}(A)$-embeddings $g_1 \colon B_1 \to C$ and $g_2 \colon B_2 \to C$ with $C \models T(A)$. Since $C \models T(A)$, it comes with an embedding $h \colon A \to C$, and we have $g_1 \circ f_1 = h = g_2 \circ f_2$. The converse is similar.

(3) $T(A)$ and $T^*(A)$ are companions. Let $B \models T(A)$. Then $B \models T$, and we have an embedding $f \colon A \to B$. Since T and T^* are companions, there is an embedding $g \colon B \to C$ with $C \models T^*$. Letting $h = g \circ f \colon A \to C$, we can view C as a model of $T^*(A)$. Since $g \circ f = h$, g is an $\mathcal{L}(A)$-embedding. The other direction is similar.

(4) $T^*(A)$ is model complete. If $g \colon M \to M'$ is an $\mathcal{L}(A)$ embedding between models of $T^*(A)$, then it is also an $\mathcal{L}$-embedding between models of T^*. Since T^* is model-complete, g is an elementary embedding in $\mathcal{L}$. Thus is preserves all first-order formulas with any parameters from M, and in particular, it preserves formulas using parameters from the image of A in M. So g is an elementary embedding in $\mathcal{L}(A)$.

Now A is an amalgamation base for T if and only if $T(A)$ has JEP. By points (3) and (4) above, $T^*(A)$ is a model companion of $T(A)$. Thus, by Theorem 6.48) $T(A)$ has JEP if and only if $T^*(A)$ is complete. □

Corollary 6.50. *Suppose T^* is a model companion of T. Then T^* is a model completion of T if and only if T has AP.*

Proof. By Theorem 6.49, T^* is a model completion of T if and only if $T^* \cup \mathrm{Diag}(A)$ is complete for all $A \models T$ if and only if every model of T is an amalgamation base if and only if T has AP. □

Corollary 6.51. *Let T be a model complete theory. T is complete if and only if $T_\forall$ has JEP, and T has QE if and only if $T_\forall$ has AP.*

Proof. Note that T and $T_\forall$ are companions, so since T is model complete, T is the model companion of $T_\forall$. Then the first assertion is direct from Theorem 6.48.

For the second assertion, by Theorem 6.7, T has QE if and only if for every substructure $A \subseteq M$ with $M \models T$, $T \cup \mathrm{Diag}(A)$ is a complete $\mathcal{L}(A)$-theory. Since the substructures of models of T are exactly the models of $T_\forall$ (Lemma 6.28), T has QE if and only if T is a model completion of $T_\forall$, if and only if $T_\forall$ has AP by Corollary 6.50. $\square$

Suppose T is inductive and has a model companion T^*. If $T_\forall$ has AP, then T^* has QE by Corollary 6.51. In this case, in the statement of Theorem 6.43, for every $\exists$-formula $\varphi(x)$, we can take the formula $\widetilde{\varphi}(x)$ to be quantifier-free. Then we can rewrite the sentence $\forall x\,(\varphi(x) \vee \widetilde{\varphi}(x))$ as

$$\forall x\,(\neg\widetilde{\varphi}(x) \to \varphi(x)),$$

or as the $\exists$-sequent

$$\widetilde{\varphi}(x) \vdash_x \varphi(x)$$

Axioms of this form are often called **extension axioms**: If some $a \in M^x$ satisfies the quantifier-free condition $\neg\widetilde{\varphi}(x)$, then it can be extended in the way described by the $\exists$-formula $\varphi(x)$.

The result of Theorem 6.43 then says that T^* can be axiomatized by these extension axioms:

$$T \cup \{\neg\widetilde{\varphi}(x) \vdash \varphi(x) \mid \varphi(x) \text{ an } \exists\text{-formula}\}.$$

7 Countable models

Lecture 21: 12/1

Throughout this rest of these lecture notes, we assume that $\mathcal{L}$ is a countable language, so that the number of $\mathcal{L}$-formulas in any finite context is countably infinite. Further, we assume that T is a complete, theory with infinite models (from which it follows by the Downward Löwenheim–Skolem Theorem, Theorem 5.13, that T has models of cardinality $\aleph_0$).

We will study the countable models of T and elementary embeddings between them. The main question we want to address is whether T has countable models which are *largest* or *smallest* with respect to elementary embeddings. More precisely, we ask whether T has a **universal** countable model, into which every other countable model embeds elementarily, or whether T have a **prime** model, which embeds elementarily into every other model. Note that prime and universal models are typically not terminal and initial objects in the category of countable models: the elementary embeddings will rarely be unique.

We will characterize the existence of universal and prime models, as well as $\aleph_0$-categoricity (the existence of a unique countable model up to isomorphism), in terms of properties of the spaces of complete types relative to T, and we will identify special kinds of countable models (**saturated** and **atomic** models, which are always universal and prime, respectively) in terms of the types they realize. We will show that the existence of a universal model implies the existence of a prime model. And we will end with a discussion of the possible numbers of countable models up to isomorphism a theory may have.

7.1 Type spaces

Recall that for a context x, a complete type in context x (relative to T) is a set $p(x)$ of formulas in context x which is satisfiable in a model of T and such that for any formula $\varphi(x)$, either $\varphi(x) \in p$ or $\neg\varphi(x) \in p$. We write $S_x(T)$ for the set of all complete types in context x (relative to T). When x is the context $\{x_1, \ldots, x_n\}$, we write $S_n(T)$ for $S_x(T)$.

Similarly, if $A \subseteq M \models T$, we write $S_x(A)$ for the set of complete types in context x with parameters from A. This is just $S_x(T(A))$, where $T(A)$ is the $\mathcal{L}(A)$-theory $\mathrm{Th}(M(A))$.

For any formula $\varphi(x)$, we define $[\varphi(x)] = \{p \in S_x(T) \mid \varphi(x) \in p\}$.

Lemma 7.1. *For any formulas $\varphi(x)$ and $\psi(x)$, φ and ψ are T-equivalent if and only if $[\varphi] = [\psi]$.*

Proof. Suppose φ and ψ are T-equivalent. By symmetry, it suffices to show that $[\varphi] \subseteq [\psi]$. So suppose $p \in [\varphi]$, i.e., $\varphi \in p$. Since φ and ψ are T-equivalent, $\{\varphi, \neg\psi\}$ is not satisfiable in any model of T, so $\neg\psi \notin p$. Since p is complete, $\psi \in p$, so $p \in [\psi]$.

Conversely, suppose φ and ψ are not T-equivalent. Without loss of generality, there is some $M \models T$ and $a \in M^x$ such that $M \models \varphi(a)$ and $M \models \neg\psi(a)$. Then $\mathrm{tp}(a) \in [\varphi] \setminus [\psi]$, so $[\varphi] \neq [\psi$. $\square$

Lemma 7.2. *Let any formula $\varphi(x)$ and $\psi(x)$,*

(a) $[\varphi(x) \wedge \psi(x)] = [\varphi(x)] \cap [\psi(x)]$.

(b) $[\varphi(x) \vee \psi(x)] = [\varphi(x)] \cup [\psi(x)]$.

(c) $[\neg\varphi(x)] = S_x(T) \setminus [\varphi(x)]$.

Proof. For (a), suppose $p \in [\varphi \wedge \psi]$, so $\varphi \wedge \psi \in p$. Since p is satisfiable, $\neg\varphi \notin p$, so since p is complete, $\varphi \in p$. Similarly, $\psi \in p$. Thus $p \in [\varphi] \cap [\psi]$. Conversely, suppose $p \in [\varphi] \cap [\psi]$, so $\varphi \in p$ and $\psi \in p$. Since p is satisfiable, $\neg(\varphi \wedge \psi) \notin p$, so since p is complete, $\varphi \wedge \psi \in p$. Thus $p \in [\varphi \wedge \psi]$.

For (c), every complete type contains either φ or $\neg\varphi$, and no complete type contains both, since complete types are satisfiable.

(b) follows from (a) and (c). □

We endow each type space $S_x(T)$ with a topology, taking as a basis of open sets

$$\{[\varphi(x)] \mid \varphi(x) \text{ a formula}\}.$$

Notice that each basic open set $[\varphi(x)]$ is closed as well (i.e., it is clopen) since it complement $[\neg\varphi(x)]$ is also open.

Theorem 7.3. *Each type space $S_x(T)$ is a **Stone space**: A zero-dimensional compact Hausdorff space.*

Proof. Zero-dimensional means that the space has a basis of clopen sets, which we have already observed.

The Hausdorff property means that for all $p \neq q$ in $S_x(T)$, there are open neighborhoods $p \in U$ and $q \in V$ with $U \cap V = \varnothing$. Since p and q are distinct complete types, there is some formula $\varphi(x)$ such that $\varphi(x) \in p$ and $\neg\varphi(x) \in q$. Then $[\varphi(x)]$ and $[\neg\varphi(x)]$ are disjoint open neighborhoods of p and q respectively.

It suffices to check compactness on open covers by basic open sets. So suppose $\{[\varphi(x)] \mid \varphi \in I\}$ is a basic open cover of $S_x(T)$. We would like to show it has a finite subcover. Since $\bigcup_{\varphi \in I}[\varphi(x)] = S_x(T)$, taking complements we have $\bigcap_{\varphi \in I}[\neg\varphi(x)] = \varnothing$, i.e., no complete type contains

$$\Sigma(x) = \{\neg\varphi(x) \mid \varphi \in I\}.$$

Therefore this set is not satisfiable (if it were satisfiable, the complete type of any realization would contain $\Sigma(x)$). By compactness, there is a finite subset $J \subseteq_{\text{fin}} I$ such that $\Sigma_0(x) = \{\neg\varphi(x) \mid \varphi \in J\}$ is unsatisfiable. Then no complete type contains $\Sigma_0(x)$, so $\bigcap_{\varphi \in J}[\neg\varphi(x)] = \varnothing$. Taking complements again, $\bigcup_{\varphi \in J}[\varphi(x)] = S_x(T)$. Thus $\{[\varphi(x)] \mid \varphi \in J\}$ is a finite subcover, and $S_x(T)$ is compact. □

7.2 Countable saturated models

Definition 7.4. Let $M \models T$ be a countable model. M is **saturated** if for all finite $A \subseteq_{\text{fin}} M$, every type $p(x) \in S_1(A)$ is realized in M.

The condition that A is finite is essential: since M is infinite, the partial type $\{x \neq m \mid m \in M\}$ is consistent by compactness, hence it extends to a complete type in $S_1(M)$, which is not realized in M.

The definition of saturated has a lot in common with the definition of existentially closed. They both capture different senses of the motto "anything that could happen (in a larger model) happens already"

Given a saturated model M, if we take an elementary extension $M \preceq N$ and $b \in N \setminus M$, for any $A \subseteq_{\text{fin}} M$, the complete type $\text{tp}(b/A)$ can be realized back in M. Compare with an EC model M, where if M is a substructure of N (not necessarily elementary) and $b \in N \setminus M$ realizes some quantifier-free formula $\varphi(x)$ with parameters from M, i.e., from a finite $A \subseteq_{\text{fin}} M$, then $\varphi(x)$ can be realized back in M.

So the main differences are that in saturated models, we consider elementary embeddings and complete types, while in EC models, we consider ordinary embeddings and quantifier-free formulas. Also, the definition of EC refers to formulas in arbitrary finite variable contexts, while the definition of saturated refers to types in one free variable. As we will see, we can bootstrap the saturation condition to arbitrary (even countable) variable contexsts, while the definition of EC has no reduction to one variable in general.

Lemma 7.5. *If $A \subseteq M$ is a set, $f\colon A \to M'$ is a partial elementary map, and $p \in S_x(A)$, then*

$$f_*p = \{\varphi(x, f(a)) \mid \varphi(x, a) \in p\}$$

is a complete type in $S_x(f(A))$.

Proof. For any $\mathcal{L}(f(A))$-formula $\varphi(x, f(a))$ in context x, either $\varphi(x, a) \in p$ or $\neg\varphi(x, a) \in p$, so either $\varphi(x, f(a)) \in f_*p$ or $\neg\varphi(x, f(a)) \in f_*p$.

So it only remains to check that f_*p is satisfiable. Since since p and hence also f_*p is closed under conjunction, by compactness it suffices to show that for every formula $\varphi(x, f(a)) \in f_*p$, $\varphi(x, f(a))$ is satisfiable in M'. Since p is satisfiable in M, $M \models \exists x\, \varphi(x, a)$, and since f is partial elementary, also $M' \models \exists x\, \varphi(x, f(a))$. $\square$

Theorem 7.6. *Let $M \models T$ be a countable saturated model. Suppose $N \models T$ is countable and $f\colon N \dashrightarrow M$ is a partial elementary map such that* $\text{dom}(f) =$ *$A \subseteq_{\text{fin}} N$. Then f extends to an elementary embedding $g\colon N \to M$.*

As a consequence, M is universal: For all countable $N \models T$, there exists an elementary embedding $N \to M$.

Proof. Enumerate N as $(n_i)_{i\in\omega}$. We define a chain $(g_i)_{i\in\omega}$ of partial elementary maps by recursion such that $\text{dom}(g_i) = A_i = A \cup \{n_j \mid j < i\}$. In the base case, take $g_0 = f$ with domain $A_0 = A$.

Given g_i, consider $p(x) = \text{tp}(n_i/A_i)$. Since $g_i(A_i)$ is finite and M is saturated, there is a realization m_i of $(g_i)_*p \in S_x(g_i(A_i))$ (by Lemma 7.5). Let g_{i+1} extend g_i by $g_{i+1}(n_i) = m_i$. Then g_{i+1} is partial elementary, since for any formula $\varphi(x_0, \ldots, x_i)$, $N \models \varphi(n_0, \ldots, n_i)$ iff $\varphi(n_0, \ldots, n_{i-1}, x) \in p$ iff $\varphi(g_i(n_0), \ldots, g_i(n_{i-1}), x) \in (g_i)_*p$ iff $M \models \varphi(g_{i+1}(n_0), \ldots, g_{i+1}(n_{i-1}), g_{i+1}(n_i))$.

Finally, the union $f = \bigcup_{i \in \omega} f_i$ is an elementary embedding $N \to M$.

For the last statement, if M is a countable saturated model and N is any countable model, the empty function is a partial elementary map $N \dashrightarrow M$, since T is complete. So it extends to some elementary embedding $N \to M$. $\square$

Corollary 7.7. *Let $M \models T$ be a countable saturated model. Let $A \subseteq_{\text{fin}} M$ and let x be a countable variable context (so $|x| \leq \aleph_0$). Then every $p \in S_x(A)$ is realized in M.*

Proof. Let $p \in S_x(A)$. Since p is a complete type, it is satisfiable by $b \in N^x$ in some elementary extension $M \preceq N$. By Löwenheim–Skolem, since A is finite and b is countable, we can take a countable elementary substructure $N' \preceq N$ such that $A \subseteq N'$ and $b \in (N')^x$. Then b realizes p in N'. Note that the identity map $A \to A$ is a partial elementary map $N' \dashrightarrow M$, since if $N' \models \varphi(a)$ with $a \in A^y$, then $N \models \varphi(a)$, and $M \models \varphi(a)$. By Theorem 7.6, there is an elementary embedding $g \colon N' \to M$ extending the identity map on A. Then $g(b) \in M^x$ realizes p. $\square$

Lecture 22: 12/3

The proof of Theorem 7.6 used a method called "going forth", where we define an elementary embedding one element at a time by recursion. In the next theorem, we will try to use the same idea to build an isomorphism. To do that, we need to ensure that the elementary embedding we build is surjective, so we need to go "back-and-forth", at each step adding one element of the domain and one element to the range.

Theorem 7.8. *Suppose $M \models T$ is a countable saturated model. Then:*

(1) If $M' \models T$ is another countable saturated model and $f \colon M \dashrightarrow M'$ is a partial elementary map with $\text{dom}(f) = A \subseteq_{\text{fin}} M$, then f extends to an isomorphism $g \colon M \cong M'$.

(2) M is the unique countable saturated model of T up to isomorphism.

(3) M is strongly homogeneous: For any finite context x and $a, a' \in M^x$, if $\text{tp}(a) = \text{tp}(a')$, then there is an automorphism $\sigma \in \text{Aut}(M)$ such that $\sigma(a) = a'$.

Proof. For (1), enumerate $M = (m_i)_{i \in \omega}$ and $M' = (m'_i)_{i \in \omega}$. We define a chain $(g_i)_{i \in \omega}$ of partial elementary maps by recursion such that for all i,

$$A_i = A \cup \{m_j \mid j < i\} \subseteq \text{dom}(g_i)$$

and

$$B_i = f(A) \cup \{m'_j \mid j < i\} \subseteq \text{ran}(g_i).$$

In the base case, take $g_0 = f$ with domain $A_0 = A$ and range $B_0 = f(A)$.

Given g_i, we first extend g_i to g'_i with $m_i \in \text{dom}(g'_i)$ exactly as in the proof of Theorem 7.6, using saturation of M'. Now we extend g'_i to g_{i+1} with $m'_i \in \text{ran}(g_{i+1})$.

Note that $h_i = (g'_i)^{-1}$ is a partial elementary map $M' \dashrightarrow M$. Extend h_i to h_{i+1} with $m'_i \in \mathrm{dom}(h_{i+1})$ exactly as in the proof of Theorem 7.6, this time using saturation of M. Now $g_{i+1} = h_{i+1}^{-1}$ is the required partial elementary map.

Finally, $g = \bigcup_{i\in\omega} g_i$ is an elementary map $M \to M'$ (because each m_i is in its domain) which is surjective (because each m'_i is in its range), and hence an isomorphism.

(2) follows immediately from (1), since if $M' \models T$ is another countable saturated model, then the empty function is a partial elementary map $M \dashrightarrow M'$, since T is complete, so it extends to an isomorphism $M \cong M'$.

(3) also follows immediately from (1), since if $\mathrm{tp}(a) = \mathrm{tp}(a')$, then the function mapping a to a' (according to their common enumeration by the variables x) is a partial elementary map $M \dashrightarrow M$, so it extends to an automorphism of M. □

You may wonder whether these definitions and results extend to uncountable models. They do have natural generalizations, we just need to use transfinite recursion for the proofs.

For an infinite (possibly uncountable) model M and an infinite cardinal κ:

- We say that a structure M is **κ-saturated** if for every $A \subseteq M$ with $|A| < \kappa$ and every $p \in S_1(A)$, p is realized in M. If M is infinite, M cannot be κ-saturated for any $\kappa > |M|$. So we say M is **saturated** if it is $|M|$-saturated, i.e., it has the maximal amount of saturation it can.

- We say that M is **κ-universal** if every $N \models T$ with $|N| \leq \kappa$ embeds elementarily in M, and we say that M is **universal** if it is $|M|$-universal.

- We say that M is **κ-strongly homogeneous** if for all x with $|x| < \kappa$ and $a, a' \in M^x$, if $\mathrm{tp}(a) = \mathrm{tp}(a')$, then there is an automorphism $\sigma \in \mathrm{Aut}(M)$ such that $\sigma(a) = a'$, and we say that M is **strongly homogeneous** if it is $|M|$-strongly homogeneous.

Now just as we proved above, uncountable saturated models are unique up to isomorphism (when they exist), universal, and strongly homogeneous. However, they do not always exist (even countable saturated models do not always exist, as we will see in a moment). It is a theorem that for any theory T with infinite models and any cardinal κ, T has a κ-saturated model (which may have cardinality much greater than κ). But if we want to work with κ-saturated models, some of the above breaks down. A κ-saturated model is always κ-universal, but it may not be κ-strongly homogeneous or unique up to isomorphism (essentially because we cannot enumerate its elements in order-type κ).

What is still true is that κ-saturated models are **κ-homogeneous**: For all x with $|x| < \kappa$ and $a, a' \in M^x$ and $b \in M$, if $\mathrm{tp}(a) = \mathrm{tp}(a')$, then there exists $b' \in M$ such that $\mathrm{tp}(ab) = \mathrm{tp}(a'b')$ (we only have to realize the appropriate type). This other notion of homogeneity is the reason that the automorphism version is called **strongly homogeneous**.

We have seen that saturated models are nice. Now we address the question of when countable saturated models exist.

Definition 7.9. We say that T is **small** if $|S_n(T)| \leq \aleph_0$ for all $n \in \omega$.

Theorem 7.10. *The following are equivalent:*

(1) *T is small.*

(2) *T has a countable saturated model.*

(3) *T has a countable universal model.*

Proof. We have already proved (2) $\Rightarrow$ (3) (Theorem 7.6).

For (3) $\Rightarrow$ (1), suppose T has a countable universal model M. Let $n \in \omega$. For any type $p \in S_n(T)$, p is realized by $a \in N^n$ in some countable model $N \models T$. By universality, there is an elementary embedding $f\colon N \to M$. Then $M \models p(f(a))$. So M realizes every type in $S_n(T)$. But there are only countably many n-tuples in M^n, and each realizes only one complete type, so there are only countably many types in $S_n(T)$.

For (1) $\Rightarrow$ (2), assume T is small. We must construct a saturated model. We begin with the observation that if $M \models T$ is any model and $A \subseteq_{\text{fin}} M$ is a finite subset, then $|S_x(A)| \leq \aleph_0$. Indeed, enumerating $A = \{a_1, \ldots, a_n\}$, and letting $y = (y_1, \ldots, y_n)$, there is a map $S_x(A) \to S_{xy}(T)$ given by $p(x, a_1, \ldots, a_n) \mapsto p(x, y_1, \ldots, y_n)$. Here is another way to describe this map: let $p(x) \in S_x(A)$, and let b be a realization of $p(x)$ in some elementary extension M' of M. Then $p(x) \mapsto \text{tp}_{M'}(ba_1 \ldots a_n)$. This map is injective, so $|S_x(A)| \leq |S_{xy}(T)| \leq \aleph_0$, since T is small.

Now let M_0 be any countable model of T. There are countably many finite subsets $A \subseteq_{\text{fin}} M_0$, and for each such A, the type space $S_x(A)$ is countable. So we can enumerate all of the types over finite subsets of M_0 as $(p_n)_{n\in\omega}$. We build an elementary chain $M_0 \preceq M_1 \preceq M_2$ such that each M_{n+1} realizes p_n and is countable (by applying Löweheim–Skolem). Let $N_0 = M_0$ and $N_1 = \varinjlim M_n$. Then $N_0 \preceq N_1$, so $N_1 \models T$, and as a countable direct limit of countable models, N_1 is countable. Moreover, N_1 realizes every type over every finite subset of N_0.

Repeating, we build an elementary chain $N_0 \preceq N_1 \preceq N_2 \preceq \ldots$ such that each N_{i+1} is countable and realizes all types over all finite subsets of N_i. Letting $N_\omega = \varinjlim N_i$, again $N_\omega \models T$ and N_ω is countable. Finally, N_ω is saturated, since if $A \subseteq_{\text{fin}} N_\omega$ and $p \in S_x(A)$, we have $A \subseteq N_k$ for some $k \in \omega$, and p is realized in N_{k+1}. □

Note that it is not true in general that every universal countable model of T is saturated. The fact is just that the *existence* of a universal countable model is equivalent to the *existence* of a saturated countable model.

As an example, $M = \mathbb{Z} \times \mathbb{Q}$ with the lexicographic order is a saturated model of the theory of discrete linear orders without endpoints, hence it is universal. Now M embeds in $N = M + \mathbb{Z}$ (appending an additional copy of $\mathbb{Z}$

whose elements are greater than all elements of M). Then N is still universal, since every countable model embeds in M, which embeds in N. But N is not saturated, since the complete type asserting that $x > 0$, $x > S(0)$, $x > S(S(0))$, etc. is not realized in N.

The conditions for the existence of saturated models of uncountable cardinality κ are much more delicate: we either need strong set-theoretic assumptions about κ (like being strongly inaccessible or an instance of GCH: $\kappa = \lambda^+ = 2^\lambda$), or strong conditions on T bounding the sizes of its type spaces over infinite sets of parameters (stability).

7.3 Model theoretic forcing

Lecture 23: 12/8

In this section, we retain our convention that $\mathcal{L}$ is countable, but we relax our focus on complete theories for a moment.

Fix a $\forall$-theory T (a set of $\forall$-sentences). We introduce a new method of building countable models of T, called "model theoretic forcing", which gives us much better control on the constructed model than black box applications of compactness.

Let $\mathcal{L}' = \mathcal{L} \cup \{c_i \mid i \in \omega\}$, where the c_i are new constant symbols. A **condition** is a finite set C of literal $\mathcal{L}'$-sentences such that $T \cup C$ is satisfiable. Let $\mathcal{C}$ be the set of conditions, partially ordered by $\subseteq$.

An **ideal** in $(\mathcal{C}, \subseteq)$ is a set $\mathcal{I} \subseteq \mathcal{C}$ of conditions with the following properties:

1. $\mathcal{I}$ is non-empty.
2. $\mathcal{I}$ is downwards closed: If $C \in \mathcal{I}$ and $C' \subseteq C$, then $C' \in \mathcal{I}$.
3. $\mathcal{I}$ is directed: If $C, C' \in \mathcal{I}$, then there exists $D \in \mathcal{I}$ with $C, C' \subseteq D$.

Using condition 2, we could replace condition 1 by $\varnothing \in \mathcal{I}$, and we could replace condition 3 by $C \cup C' \in \mathcal{I}$.

Given an ideal $\mathcal{I}$, $\Delta_{\mathcal{I}} = \bigcup \mathcal{I}$ is a set of literal $\mathcal{L}'$-sentences. Since every finite subset of $\Delta_{\mathcal{I}}$ is a condition $C \in \mathcal{I}$ (using conditions 2 and 3), it follows by compactness that $T \cup \Delta_{\mathcal{I}}$ is satisfiable. Conversely, if Δ is a set of literal $\mathcal{L}'$-sentences such that $T \cup \Delta$ is satisfiable, it is easy to check that $\downarrow\Delta = \mathcal{P}_{\text{fin}}(\Delta)$ is an ideal in $\mathcal{C}$.

The main idea of model theoretic forcing is this: We want to carefully build an ideal $\mathcal{I}$ by assembling it from conditions in $\mathcal{C}$ such that $\Delta_{\mathcal{I}}$ is the diagram of a model of T with the properties we want. The properties we can ensure through this construction are called "enforceable".

Let U be a set of conditions. We say U is **open** if it is upwards closed: if $C \in U$ and $C \subseteq C'$, then $C' \in U$. We say U is **dense** if for every $C \in \mathcal{C}$, there exists $D \in U$ with $C \subseteq D$. The reason for the topological names is that we can put a topology on $\mathcal{C}$ whose basic open sets have the form $U_C = \{D \in \mathcal{C} \mid C \subseteq D\}$ for $C \in \mathcal{C}$. In this topology, dense and open have their usual meanings. However, we won't use this topology in any way other than this terminology.

Let P be a set of ideals. We say that P is **enforceable** if there is a countable set $\mathcal{U}$ of dense open sets in $\mathcal{C}$ such that for any ideal $\mathcal{I}$, if $\mathcal{I} \cap U \neq \varnothing$ for all $U \in \mathcal{U}$, then $\mathcal{I} \in P$.

Theorem 7.11 (Rasiowa–Sikorski lemma)**.** *If P is an enforceable set of ideals, then for any condition C, there exists an ideal $\mathcal{I} \in P$ with $C \in \mathcal{I}$.*

Proof. Enforceability of P is witnessed by $\mathcal{U} = \{U_i \mid i \in \omega\}$ where each U_i is dense and open. We build a sequence of conditions $(C_i)_{i\in\omega}$ by recursion. Let $C_0 = C$. Given C_i, by density of U_i, we can pick $C_{i+1} \in U_i$ with $C_i \subseteq C_{i+1}$.

Now let $\mathcal{I} = \downarrow\bigcup_{i\in\omega} C_i = \{D \in \mathcal{C} \mid D \subseteq C_i \text{ for some } i \in \omega\}$. This is an ideal in $\mathcal{C}$. We have $C = C_0 \in \mathcal{I}$ and for all $i \in \omega$, $C_{i+1} \in \mathcal{I} \cap U_i$, so $\mathcal{I} \in P$. □

Here is another way of looking at the construction that might better motivate the definition of enforceability. Consider a game between two players, ∀belard and ∃loise. ∀belard goes first, picking a condition C_0. The players then take turns picking conditions C_i such that $C_i \subseteq C_{i+1}$ for all i. After playing forever, we let $\mathcal{I} = \downarrow\bigcup_{i\in I} C_i$. If P is enforceable, then ∃loise has a strategy to "force" $\mathcal{I} \in P$: On her i^{th} turn, she picks condition C_{2i+1} such that $C_{2i} \subseteq C_{2i+1}$ and $C_{2i+1} \in U_i$.

The key property of the family of enforceable sets is that it is closed under countable intersections.

Theorem 7.12. *If $(P_n)_{n\in\omega}$ is a countable family of sets of ideals, each of which is enforceable, then $\bigcap_{n\in\omega} P_n$ is enforceable.*

Proof. For each $n \in \omega$, pick $\mathcal{U}_n$ witnessing enforceability of P. Then $\mathcal{U}_\omega = \bigcup_{n\in\omega} \mathcal{U}_n$ is a countable set of dense open sets in $\mathcal{C}$. For any ideal $\mathcal{I}$, if $\mathcal{I} \cap U \neq \varnothing$ for all $U \in \mathcal{U}_\omega$, then $\mathcal{I} \in P_n$ for all $n \in \omega$, so $\mathcal{I} \in \bigcap_{n\in\omega} P_n$. Thus $\mathcal{U}_\omega$ witnesses enforceability of $\bigcap_{n\in\omega} P_n$. □

Let's now check enforceability of some desirable properties. For a condition C, let $\text{supp}(C)$ be the finite set of constant symbols mentioned in sentences in C.

Lemma 7.13. *The set of ideals $\mathcal{I}$ such that $\Delta_{\mathcal{I}}$ is complete (either φ or $\neg\varphi$ is in $\Delta_{\mathcal{I}}$ for each literal $\mathcal{L}'$-sentence φ) is enforceable.*

Proof. For each literal $\mathcal{L}'$-sentence φ, let $U_\varphi = \{C \mid \varphi \in C \text{ or } \neg\varphi \in C\}$. Then U_φ is clearly open. For density, let C be a condition, and let $M \models T \cup C$. Then $M \models \varphi$ or $M \models \neg\varphi$, so at least one of $C \cup \{\varphi\}$ and $C \cup \{\neg\varphi\}$ is a condition in U_φ. If $\mathcal{I}$ meets U_φ for all φ, then $\Delta_{\mathcal{I}}$ is complete, so this set of ideals is enforceable. □

Lemma 7.14. *The set of ideals $\mathcal{I}$ such that for all closed $\mathcal{L}'$-terms t there exists $i \in \omega$ such that $t = c_i \in \Delta_{\mathcal{I}}$ is enforceable.*

Proof. For each closed $\mathcal{L}'$-term t, let $U_t = \{C \mid \exists i \in \omega, t = c_i \in C\}$. Then U is clearly open. For density, let C be a condition. Pick any $i \in \omega \setminus \text{supp}(C)$. Then $C \cup \{t = c_i\}$ is a condition in U_t. If $\mathcal{I}$ meets U_t for all t, then $\Delta_{\mathcal{I}}$ satisfies the property in the statement, so this set of ideals is enforceable. □

Suppose $\mathcal{I}$ is an ideal such that $\Delta_\mathcal{I}$ is complete and for all closed $\mathcal{L}'$-terms t there exists $i \in \omega$ such that $t = c_i \in \Delta_\mathcal{I}$. Let $N \models T \cup \Delta_\mathcal{I}$. Then the subset $\{c_i^N \mid i \in \omega\}$ is $\mathcal{L}'$-closed, so it is the domain of an $\mathcal{L}'$-substructure $M_\mathcal{I} \subseteq N$. Since $T \cup \Delta_\mathcal{I}$ is a $\forall$-theory, $M \models T \cup \Delta_\mathcal{I}$. And since $\Delta_\mathcal{I}$ is complete, $\Delta_\mathcal{I} = \mathrm{Diag}(M_\mathcal{I})$. We call $M_\mathcal{I}$ the **compiled model** corresponding to $\mathcal{I}$.

We have shown (Lemmas 7.13 and 7.14) that it is enforceable that $\mathcal{I}$ has a compiled model. Given a property P of countable structures, we can now write "P is enforceable" as shorthand for "the set of all ideals $\mathcal{I}$ such $\mathcal{I}$ has a compiled model $M_\mathcal{I}$ and $M_\mathcal{I}$ satisfies P is enforceable".

Theorem 7.15. *Let T be an inductive theory. In model-theoretic forcing with respect to $T_\forall$, it is enforceable that the compiled model is an EC model of T.*

Proof. By Lemma 6.36, every EC model of $T_\forall$ is an EC model of T. So it suffices to show that it is enforceable that the compiled model (which is automatically a model of $T_\forall$) is an EC model of $T_\forall$.

Let $\varphi(x)$ be an $\exists$-formula and let c be an assignment of x to the constant symbols. Recall the notation Cont_φ^T from Lemma 6.42. Define

$$U_{\varphi(c)} = \{C \mid T_\forall \cup C \models \varphi(c) \text{ or there is } \psi(x,y) \in \mathrm{Cont}_\varphi^{T_\forall} \text{ and constants } d \text{ such that } T_\forall \cup C \models \psi(c,d)\}.$$

Then $U_{\varphi(c)}$ is clearly open. For density, let C be any condition, and suppose $M \models T_\forall \cup C$. Embed M in an EC model $N \models T_\forall \cup C$. Since N is EC, by Lemma 6.42 either $N \models \varphi(c)$ or $N \models \psi(c,e)$ for some $e \in N^y$ and $\psi(x,y) \in \mathrm{Cont}_\varphi^{T_\forall}$.

If $N \models \varphi(c)$, write $\varphi(x)$ in normal form as $\bigvee_{i=1}^k \exists y^i \varphi_i(x,y^i)$, where φ_i is a conjunction of literals. Then of the disjuncts $\exists y^i \varphi_i(c,y^i)$ is true in N. Let $b \in N^{y_i}$ be witnesses to the quantifiers, so $N \models \varphi_i(c,b)$. Now pick c' to be an interpretation of y^i in the constant symbols outside of $\mathrm{supp}(C)$, and let $C' = C \cup \{\chi(c,c') \mid \chi \text{ a conjunct of } \varphi_i\}$. Then C' is a condition in $U_{\varphi(c)}$, since $T_\forall \cup C' \models \varphi(c)$.

If $N \models \psi(c,e)$ with $\psi \in \mathrm{Cont}_\varphi^{T_\forall}$, we can use similar reasoning to build a condition in $U_{\varphi(c)}$ extending C.

Now if $\mathcal{I}$ is an ideal meeting $U_{\varphi(c)}$ for all $\varphi(c)$, the compiled model $M_\mathcal{I}$ will have the property that for any $\exists$-formula $\varphi(x)$ and any $c \in M^x$, either $M \models \varphi(c)$ or $M \models \psi(c,d)$ for some $\psi(x,y) \in \mathrm{Cont}_\varphi^{T_\forall}$ and some constants d. Thus $M_\mathcal{I}$ will be an EC model of $T_\forall$ by Lemma 6.42. □

I'll end this general discussion of model theoretic forcing with a comment about the connection to set theoretic forcing. Set theoretic forcing starts with a poset $(P, \leq)$ whose elements are called conditions. This generalizes our $(\mathcal{C}, \subseteq)$, although the usual convention in forcing is that the poset is upside-down: if $p \leq q$, then p is the "stronger" condition. As a result, set theorists talk about filters on P, rather than ideals.

A filter is called **generic** if it meets *every* dense open set in the poset. The Rasiowa–Sikorski lemma shows that for any countable collection of dense open sets, there is a filter meeting all of them, but there may be no generic filter. However, set theoretic forcing starts with a model V of set theory and builds a new universe of set theory $V[G]$ containing a V-generic filter G on P, i.e., one which meets every dense open set present in V.

If T is an inductive theory and we carry out set theoretic forcing with our poset $(\mathcal{C}, \subseteq)$ relative to $T_\forall$, we have shown that $V[G]$ will contain a compiled model M which is an EC model of T (and much more). This "generic model" may or may not be isomorphic to a structure in V, but we can view its properties as "generic properties" of models of T. For this reason, properties of EC models of T (e.g., the axioms of the model companion, if it exists) are often called "generic" properties relative to T.

7.4 Omitting types and prime and atomic models

Lecture 24: 12/10

In this section, we will use the machinery of model-theoretic forcing to build models that omit (i.e., fail to realize) certain types. We return to our setting of a complete theory T with infinite models in a countable language $\mathcal{L}$.

Let $p(x) \in S_x(T)$ be a complete type. We say that p is **isolated** if there is a formula $\varphi(x)$ such that $\varphi(x) \in p(x)$, and for every formula $\psi(x) \in p(x)$, $T \models \forall x\,(\varphi(x) \to \psi(x))$. In other words, $\varphi(x)$ completely determines the other formulas in $p(x)$, so $p(x)$ is the only complete type in $S_x(T)$ containing $\varphi(x)$.

The reason for the name is that p is an isolated point in the topology on $S_x(T)$. The basic clopen set $[\varphi] = \{q \in S_x(T) \mid \varphi \in q\}$ contains only p.

Example 7.16. An example of an isolated type relative to ACF_0 is $\mathrm{tp}(i/\varnothing)$, which is isolated by the formula $x^2 + 1 = 0$. In general, for any irreducible polynomial $p \in \mathbb{Q}[x]$, $p(x) = 0$ isolates a type.

For another example, if T is DLO, there is only one type in $S_1(T)$, so this type is isolated (by $x = x$). There are three types in $S_2(T)$, isolated by the formulas $x_1 < x_2$, $x_1 = x_2$, and $x_1 > x_2$.

Proposition 7.17. *Let $p \in S_x(T)$ be an isolated type. Then p is realized in every model of T.*

Proof. Suppose p is isolated by $\varphi(x)$, and let $M \models T$ be any model. Since p is satisfiable, we can pick some realization $b \in N^x$ of p in $N \models T$. Then $N \models \varphi(b)$, so $N \models \exists x\,\varphi(x)$, and since T is complete, $M \models \exists x\,\varphi(x)$. Letting $a \in M^x$ be a witness, a realizes p in M. ☐

Conversely, we show that if $p \in S_x(T)$ is *not* isolated, then it can be omitted.

Theorem 7.18 (Omitting Types)**.** *Let $p \in S_n(T)$ be a non-isolated type. Then there exists a countable model $M \models T$ such that p is not realized in M.*

Proof. Let $\mathcal{L}_{\mathrm{FO}^*}$ be the Morleyized language constructed in Section 4.4, which has an n-ary relation symbol R_φ for every formula $\varphi(x_1, \ldots, x_n)$. Let $\widehat{T}$ be the $\exists^+$-theory which is a definable expansion of T. We consider model-theoretic forcing with respect to $\widehat{T}_\forall$. By Theorem 7.15, it is enforceable that the ideal $\mathcal{I}$ has compiled model $M_\mathcal{I}$ with $M_\mathcal{I} \models \widehat{T}$, so $M_\mathcal{I}|_\mathcal{L} \models T$.

To show that p is not realized in $M_\mathcal{I}|_\mathcal{L}$, it suffices by Theorem 7.12 to show that for each of the countably many n-tuples of constants c, it is enforceable that p is not realized on c.

Let $U = \{C \mid \text{there is } \psi(x) \notin p(x) \text{ s.t. } R_\psi(c) \in C\}$. U is clearly open. For density, let C be a condition. The conjunction of the literal sentences in C is equivalent to an $\mathcal{L}$-formula $\varphi(c, d)$ (where d is a tuple of additional constant symbols from $\operatorname{supp}(C)$). Now the formula $\exists y\, \varphi(x, y)$ does not isolate $p(x)$, so there is a formula $\psi(x) \notin p(x)$ such that $(\exists y\, \varphi(x, y)) \wedge \psi(x)$ is satisfiable. Let $C' = C \cup \{R_\psi(c)\}$. Then C' is a condition in U.

If $\mathcal{I}$ meets U, then in $M_\mathcal{I}|_\mathcal{L}$, $c^{M_\mathcal{I}}$ satisfies some $\psi(x) \notin p(x)$, so c does not realize $p(x)$. □

An atomic model is one that only realizes the types that can't be omitted.

Definition 7.19. A model $M \models T$ is **atomic** if for all $n \in \omega$, every type in $S_n(T)$ realized in M is isolated.

Recall that $M \models T$ is prime if for every $N \models T$, there is an elementary embedding $M \to N$.

Corollary 7.20. *Suppose M is a prime model of T. Then M is countable and atomic.*

Proof. By Löwenheim–Skolem, T has a countable model N. Since M embeds elementarily in N, M is countable.

Now let $p \in S_n(T)$ and assume $a \in M^x$ realizes p. Assume for contradiction that p is not isolated. By Theorem 7.18, there is a countable model $N \models T$ such that p is not realized in N. Let $f \colon M \to N$ be an elementary embedding. Then $f(a)$ realizes p in N, contradiction. □

If we seek a prime model of T, it must be atomic. To build an atomic model, we need to omit not just one type, but *all* non-isolated types at once.

Under what conditions can we omit a set of types in a single model? it turns out that the correct analog of a type being non-isolated is a set of types being nowhere dense.

Definition 7.21. Let S be a topological space. A set $X \subseteq S$ is **nowhere dense** if for every non-empty open set $U \subseteq S$, $X \cap U$ is not dense in U. That is, for every open $U \subseteq S$, there exists a non-empty open $V \subseteq U$ such that $X \cap V = \varnothing$.

You should think of nowhere dense sets as being "very small" in a topological sense. There is a strong analogy between nowhere dense sense in topology and null sets in measure theory. Recall that a countable union of sets of measure 0 still has measure 0. It is not true that a countable union of nowhere dense sets is nowhere dense, but this suggests the following definition.

Definition 7.22. Let S be a topological space. A set $X \subseteq S$ is **meager** if it can be written as a countable union of nowhere dense sets.

Note that if p is an isolated point, then p is dense in the open set $\{p\}$. On the other hand, suppose we are in a T_1 space, so every point is closed (every Hausdorff space is T_1, so this applies to our Stone spaces of types). I claim that if p is non-isolated, then $\{p\}$ is nowhere dense. Indeed, if U is non-empty and open, then $V = U \setminus \{p\}$ is non-empty and open and disjoint from $\{p\}$. It follows that any countable set of non-isolated points is meager.

For example, although $\mathbb{Q}$ is dense in $\mathbb{R}$, it is meager, since $\mathbb{Q} = \bigcup_{q\in\mathbb{Q}}\{q\}$, a countable union of nowhere dense singleton sets.

Theorem 7.23 (Improved Omitting Types). *For each $n \in \omega$, let $X_n \subseteq S_n(T)$ be a meager set. Then there exists a countable model $M \models T$ such that for all $n \in \omega$ and for all $p \in X_n$, p is not realized in M.*

Proof. We modify the proof of Theorem 7.18. Fix $n \in \omega$ and write $X_n = \bigcup \mathcal{N}$, where $\mathcal{N}$ is a countable family of nowhere dense sets. Let $Y \in \mathcal{N}$, and let c be an n-tuple of constants. I claim it is enforceable that $\mathrm{tp}_{\mathcal{L}}(c) \notin Y$.

Let $U = \{C \mid R_\psi(c) \in C$ for some $\psi(x)$ such that for all $p(x) \in Y, \psi \notin p\}$. U is clearly open. For density, let C be a condition. The conjunction of the literal sentences in C is equivalent to an $\mathcal{L}$-formula $\varphi(c, d)$ (where d is a tuple of additional constant symbols from $\mathrm{supp}(C)$). Consider the open set $[\exists y\, \varphi(x, y)]$ in $S_n(T)$. Since C is satisfiable, this set is non-empty. Since Y is nowhere dense, there exists a non-empty open set $V \subseteq [\exists y\, \varphi(x, y)]$ such that $Y \cap V = \varnothing$. Shrinking V, we may assume it is a basic open set $[\psi(x)]$. Thus $[\psi(x)] \cap Y = \varnothing$, so $\psi(x) \notin p(x)$ for all $p \in Y$. Since $[\psi(x)]$ is non-empty, $\psi(x)$ is satisfiable, so $C' = C \cup \{R_\psi(c)\}$ is a condition in U.

If $\mathcal{I}$ meets U, then in $M_{\mathcal{I}}|_{\mathcal{L}}$, $c^{M_{\mathcal{I}}}$ satisfies some $\psi(x)$ such that $\psi \notin p$ for all $p \in Y$, so $\mathrm{tp}_{\mathcal{L}}(c) \notin Y$.

Since there are countably many tuples c, it is enforceable that $M_{\mathcal{I}}$ does not realize any type in Y. Since there are countably many $Y \in \mathcal{N}$, it is enforceable that $M_{\mathcal{I}}$ does not realize any type in X_n. Since there are countably many n, it is enforceable that no type in any X_n is realized in $M_{\mathcal{I}}$. ☐

We say that **isolated types are dense** relative to T if for all $n \in \omega$, the set of isolated types is dense in the space $S_n(T)$. Translating, this means that for every non-empty basic open set $[\varphi(x)]$, there is an isolated type $p(x) \in [\varphi(x)]$, i.e., every satisfiable formula is contained in some isolated type.

Theorem 7.24. *T has a countable atomic model if and only if isolated types are dense relative to T.*

Proof. Suppose T has a countable atomic model M. Let $\varphi(x)$ be a satisfiable formula. By completeness, $T \models \exists x\, \varphi(x)$, so there is some $a \in M^x$ such that $M \models \varphi(a)$. But then $\varphi(x) \in \mathrm{tp}(a)$, which is isolated, since M is atomic. Thus isolated types are dense.

Conversely, suppose isolated types are dense. I claim that for all $n \in \omega$, then set $\mathcal{N}$ of non-isolated types in $S_n(T)$ is nowhere dense (and hence meager). Indeed, for any non-empty open set $U \subseteq S_n(T)$, we can find an isolated type

$p \in U$, so $\{p\}$ is open and $\mathcal{N} \cap \{p\} = \varnothing$. By Theorem 7.23, there is a countable model $M \models T$ which does not realize any non-isolated type, hence is atomic. $\square$

We showed earlier that a prime model must be countable and atomic. We now show that the converse is true, and that just like countable saturated models, countable atomic models are strongly homogeneous and unique up to isomorphism.

Theorem 7.25. *Suppose $M \models T$ is a countable atomic model. Then:*

(1) M is prime.

(2) M is unique up to isomorphism.

(3) M is strongly homogeneous: If $a, a' \in M^x$ such that $\mathrm{tp}(a) = \mathrm{tp}(a')$, then there is an automorphism $f\colon M \to M$ such that $f(a) = a'$.

Proof. For (1), we go forth. Let $N \models T$. Enumerate M as $(m_i)_{i\in\omega}$. Let f_0 be the empty function, which is partial elementary $M \dashrightarrow N$ because T is complete. We define a partial elementary map f_i with domain $A_i = \{m_0, \ldots, m_{i-1}\}$ for each $i \in \omega$ by recursion. Given f_i, consider $p = \mathrm{tp}(m_0, \ldots, m_i)$. Since M is atomic, $p(x)$ is isolated by some formula $\varphi(x_0, \ldots, x_i)$.

Since $M \models \exists x_i \varphi(m_0, \ldots, m_{i-1}, x_i)$ and f_i is partial elementary, also $N \models \exists x_i \varphi(f_i(m_0), \ldots, f_i(m_{i-1}), x_i)$. Define $f_{i+1}(m_i)$ to be any witness, so $N \models \varphi(f_{i+1}(m_0), \ldots, f_{i+1}(m_i))$. Since φ isolates the complete type p relative to T, $(f_{i+1}(m_0), \ldots, f_{i+1}(m_i))$ satisfies p, so f_{i+1} is partial elementary.

The union $f = \bigcup_{i\in\omega} f_i$ is an elementary embedding $M \to N$.

For (2) and (3), we only have to modify the proof of (1) to go back and forth, just as in the proof of Theorem 7.8. $\square$

Lecture 25: 12/15

To recap, we have proved the following result, characterizing the existence of prime models by a condition on the type space.

Corollary 7.26. *The following are equivalent, for a complete theory T in a countable language:*

(1) Isolated types are dense.

(2) T has a countable atomic model.

(3) T has a prime model.

Proof. Corollary 7.20, Theorem 7.24, and Theorem 7.25. $\square$

Note that we also proved something stronger: If $M \models T$ is countable, then M is prime if and only if M is atomic, and this M is unique up to isomorphism. This is better than the situation with universal and saturated models: Every countable saturated model is universal, countable saturated models are unique up to isomorphism if they exist, and the existence of a countable universal model implies the existence of a countable saturated model, but T may have countable universal models which are not saturated.

All of our results on prime and atomic models used strongly the assumption $|\mathcal{L}| \leq \aleph_0$ (which was necessary for the omitting types theorem). In an uncountable language, a theory may have a prime model which is not atomic or an atomic model which is not prime, and prime and atomic models need not be unique up to isomorphism.

Corollary 7.27. *Suppose T has a countable universal model. Then T has a prime model.*

Proof. By Theorem 7.10, T is small. For all $n \in \omega$, let

$$N_n = \{p(x) \in S_n(T) \mid p(x) \text{ is non-isolated}\}.$$

Then since $|S_n(T)| \leq \aleph_0$, we have $|N_n| \leq \aleph_0$, so P_n is meager in $S_n(T)$ (being a countable union of non-isolated singleton sets). By Theorem 7.23, T has a countable atomic model, and hence a prime model by Corollary 7.26. □

7.5 $\aleph_0$-categorical theories

We will now characterize the theories in which prime and universal models coincide. These are the $\aleph_0$-categorical theories, which have a unique countable model up to isomorphism. This theorem is often called the Ryll-Nardzewski theorem, and additionally attributed to Engeler and Svenonius, all independently.

Theorem 7.28 (Ryll-Nardzewski)**.** *The following are equivalent:*

(1) T is $\aleph_0$-categorical, i.e., T has a unique countable model up to isomorphism.

(2) Every countable model of T is saturated.

(3) T has a countable model which is both universal and prime.

(4) Every countable model of T is atomic.

(5) For all $n \in \omega$, every type in $S_n(T)$ is isolated.

(6) For all $n \in \omega$, $S_n(T)$ is finite.

(7) For all $n \in \omega$, there are finitely many formulas in context $x = \{x_1, \ldots, x_n\}$ up to T-equivalence.

Proof. (1) ⇒ (2): Let M be the unique countable model of T. Let $n \in \omega$. Then every type in $S_n(T)$ is realized in a countable model of T, and hence in M. Thus $|S_n(T)| \leq \aleph_0$. Then T is small, so it has a countable saturated model by Theorem 7.10. By $\aleph_0$-categoricity, every countable model is saturated.

(2) ⇒ (3): Since T has a saturated countable model, it has a countable universal model, so by Corollary 7.27, T has a prime model M, which must be countable. By (2), M is also saturated, hence universal.

(3) ⇒ (4): Let $M \models T$ be countable, universal, and prime. Let $N \models T$ be a countable model. I claim that N is atomic. For any $N' \models T$ countable,

we can pick an elementary embedding $f\colon N \to M$ (since M is universal) and an elementary embedding $g\colon M \to N'$ (since M is prime), so $g \circ f\colon N \to N'$ witnesses that N is prime. Thus N is atomic by Corollary 7.20.

(4) $\Rightarrow$ (1): Since atomic models are unique up to isomorphism by Theorem 7.25, if every countable model is atomic, T is $\aleph_0$-categorical.

(4) $\Leftrightarrow$ (5): Let $n \in \omega$ and $p \in S_n(T)$. Then p is realized in some countable model $M \models T$. By (5), M is atomic. So p is isolated. Conversely, if every type is isolated, then every model realizes only isolated types, so every model is atomic.

(5) $\Leftrightarrow$ (6): By Lemma **??**, $S_n(T)$ is a compact Hausdorff space, and any compact Hausdorff space is finite if and only if it is discrete.

(6) $\Leftrightarrow$ (7): If there are k formulas in context x up to T-equivalence, then $|S_x(T)| \leq 2^k$, since a type is a set of formulas, and complete types are closed under T-equivalence (Lemma 7.1). Conversely, if φ and ψ are not T-equivalent, then by Lemma 7.1, $[\varphi(x)] \neq [\psi(x)]$ in $S_x(T)$. If $|S_x(T)| = k$, then the number of formulas in context x up to T-equivalence is bounded above by the number of subsets of $|S_x(T)|$, which is 2^k. □

Here are two example applications.

Corollary 7.29. *Suppose T is a complete theory in a finite relational language. If T has quantifier elimination, then T is $\aleph_0$-categorical.*

Proof. In a finite relational language, there are only finitely many atomic formulas in context x for any finite x. Then the number of boolean combinations of such formulas, up to equivalence, is finite (at most 2^{2^k}, when k is the number of atomic formulas). Since T has quantifier elimination, every formula in context x is equivalent to a boolean combination of atomic formulas. By Theorem 7.28, T is $\aleph_0$-categorical. □

Example 7.30. DLO is $\aleph_0$-categorical, since it has QE in a finite relational language.

The more concrete way to prove $\aleph_0$-categoricity of DLO is to use a back-and-forth argument to construct an isomorphism between any two countable models (this is due to Cantor!). Note that this is hidden in our argument: The finiteness of the type spaces implies every countable model is atomic, and we obtained $\aleph_0$-categoricity by uniquness of atomic models using back-and-forth.

Example 7.31. Let K be an infinite field. Then $T = \mathrm{Th}(K)$ is never $\aleph_0$-categorical. Indeed, by compactness we can find an elementary extension $K \preceq K'$ containing an element α which is transcendental over the prime field. It follows that the powers $\alpha, \alpha^2, \alpha^3, \ldots$ are all distinct (otherwise α would be a root of the polynomial $x^i - x^j$ for some $i \neq j$). But then the formulas $\{y = x^n \mid n \in \mathbb{N}\}$ are pairwise non-T-equivalent, since $y = x^n$ is the only one which is satisfied by $(\alpha^n, \alpha))$. By Theorem 7.28, T is not $\aleph_0$-categorical.

We have now completed our classification of complete theories with infinite models in countable languages, according to the existence of prime and universal models, in terms of the topology and cardinality of the type spaces $S_x(T)$:

- $S_n(T)$ is finite for all n ⇔ every type in $S_n(T)$ is isolated for all n: T is $\aleph_0$-categorical, and the unique countable model is both prime and universal. Example: DLO.

- $S_n(T)$ is countable or finite for all n, and $S_n(T)$ is infinite for some n: T has both a prime and a universal model, and these are different. Example: ACF_0, see Example 7.32

- Isolated types are dense, and $S_n(T)$ is uncountable for some n: T has a prime model but no universal model. Example: $\mathrm{Th}(\mathbb{Q}; <, (q)_{q\in\mathbb{Q}})$, see Example 7.33.

- Isolated types are not dense: T has neither a prime model nor a universal model. Example: $\mathrm{Th}(2^\omega; (R_n)_{n\in\omega})$, see Example 7.34.

Example 7.32. A countable algebraically closed field of characteristic 0 is determined up to isomorphism by its transcendence degree over $\mathbb{Q}$, which can be finite or countably infinite. So, up to isomorphism, the countable models of ACF_0 are

$$\overline{\mathbb{Q}} \subseteq \overline{\mathbb{Q}(t_0)} \subseteq \overline{\mathbb{Q}(t_0, t_1)} \subseteq \cdots \subseteq \overline{\mathbb{Q}(t_0, t_1, t_2, \dots)}$$

The field of countable transcendence degree is the countable saturated and universal model, while the field $\overline{\mathbb{Q}}$ of algebraic numbers is the atomic and prime model.

Example 7.33. Let $T = \mathrm{Th}(\mathbb{Q}; <, (q)_{q\in\mathbb{Q}})$, the theory of the rational order with a constant naming each element. By quantifier elimination, a type $p(x)$ in one variable x is determined by the formulas of the form $x = q$, $x < q$, and $q < x$ in $p(x)$, for $q \in \mathbb{Q}$. In particular, for every $q \in \mathbb{Q}$, there is a type $p_q(x)$ isolated by the formula $x = q$. And for every downwards-closed set $L \subseteq \mathbb{Q}$, there is a non-isolated type $p_L(x)$ which contains $\{q < x \mid q \in L\} \cup \{x < q \mid q \notin L\}$.

Since there is one cut in $\mathbb{Q}$ for every real number, T is not small, so it does not have a countable saturated model. But the isolated types (in this case, the types corresponding to the constant symbols) are dense, and T has an atomic model, namely $\mathbb{Q}$.

Example 7.34. Let $\mathcal{L} = ((R_n)_{n\in\omega})$, where each R_n is a unary relation symbol. Consider the $\mathcal{L}$-structure C with domain 2^ω, the set of all infinite binary sequences (equivalently the Cantor space), such that R_n^C is the set of all sequences such that the n^{th} term is 1. Let $T = \mathrm{Th}(C)$.

It is possible to show that T has quantifier elimination, so a complete type $p(x)$ in one variable x is determined by the set $\{n \mid R_n(x) \in p(x)\}$. It follows that $S_x(T)$ is homeomorphic to the Cantor space 2^ω. This space has no isolated types. So T has no countable saturated model and no countable atomic model.

A countable model realizes only countably many types in $S_x(T)$, and we can use the improved omitting types theorem to build a countable model omitting the types in any meager subset of the Cantor space $S_x(T)$.

7.6 The number of countable models

Given a theory T and a cardinal κ, we write $I(T, \kappa)$ for the number of models of T of cardinality κ, up to isomorphism. We will focus here on $\kappa = \aleph_0$, retaining our countable language hypothesis.

Note that every countable model is isomorphic to one with domain ω. Each n-ary relation symbol has at most $2^{\aleph_0^n} = 2^{\aleph_0}$ interpretations on ω, and each n-ary function symbol has at most $\aleph_0^{\aleph_0^n} = 2^{\aleph_0}$ interpretations on ω. So the number of countable models of T up to isomorphism is at most the number of $\mathcal{L}$-structures with domain ω, which is at most $\prod_{\mathcal{L}} 2^{\aleph_0} = (2^{\aleph_0})^{\aleph_0} = 2^{\aleph_0}$.

What are the possible values of $I(T, \aleph_0)$? Here is a summary of what we know so far:

- Any $\aleph_0$-categorical theory T has $I(T, \aleph_0) = 1$.
- It is possible to have $I(T, \aleph_0) = \aleph_0$. For example, take $T = \mathrm{ACF}_0$.
- If T is a theory with $S_n(T) = 2^{\aleph_0}$ for some n, such as the theories in Examples 7.33 and 7.34, then $I(T, \aleph_0) = 2^{\aleph_0}$. This is because every type is realized in some countable model, but any countable model can only realize countably many types.

Here is an example showing that $I(T, \aleph_0)$ can be finite but not 1.

Example 7.35. Let $T = \mathrm{Th}(\mathbb{Q}; <, (c_n)_{n\in\omega})$, where the constant symbol c_n is interpreted as the natural number n. Then T can be axiomatized by the theory of dense linear orders without endpoints, together with axioms $c_n < c_m$ when $n < m$. It is an exercise to show that these axiomatize a complete theory.

Now T has exactly 3 countable models up to isomorphism:

1. Atomic: The sequence (c_n) has no upper bound. Any such model is isomorphic to $(\mathbb{Q}; <, (c_n)_{n\in\omega})$, where $\lim_{n\to\infty} c_n = \infty$.
2. Saturated: The sequence (c_n) has an upper bound, but no least upper bound. Any such model is isomorphic to $(\mathbb{Q}; <, (c_n)_{n\in\omega})$, where $\lim_{n\to\infty} c_n = \pi$.
3. Neither atomic nor saturated: The sequence (c_n) has a least upper bound. Any such model is isomorphic to $(\mathbb{Q}; <, (c_n)_{n\in\omega})$, where $\lim_{n\to\infty} c_n = 1$.

One can modify the previous example to find complete theories T with $I(T, \aleph_0) = n$ for all $n \geq 3$. Curiously, the case $n = 2$ is impossible. It's worthwhile keeping Example 7.35 in mind while working through the proof.

Theorem 7.36 (Vaught)**.** *There is no complete theory T with exactly two countable models up to isomorphism.*

Proof. Let T be a complete theory, and assume for contradiction that T has exactly two countable models up to isomorphism. Since every type is realized in one of these two models, T is small, so T has a countable saturated model

M_1 (by Theorem 7.10) and a countable atomic model M_0 (by Corollary 7.27). Since T is not $\aleph_0$-categorical, by Theorem 7.28 there is some context $n \in \omega$ and some non-isolated type $p(x) \in S_n(T)$. This type is realized in M_1 but not in M_0, so $M_0 \not\cong M_1$. Our goal is now to find a third model of T.

Let $a \in M_1^n$ be a realization of p, and consider the expanded language $\mathcal{L}(a)$ by new constants naming the elements of a. Let $M_1(a)$ be the expansion of M_1 to a $\mathcal{L}(a)$-structure, and let $T(a) = \mathrm{Th}_{\mathcal{L}(a)}(M_1(a))$.

Now $M_1(a)$ is still a countable saturated model of $T(a)$, because any $\mathcal{L}(a)$-type over a finite set $B \subseteq M_1(a)$ is an $\mathcal{L}$-type over the finite set $B \cup a$, and hence is realized in M_1. So $T(a)$ is small, and hence it has a countable atomic model $M_{1/2}(a)$.

$T(a)$ is not $\aleph_0$-categorical by Theorem 7.28, since any infinite family of $\mathcal{L}$-formulas which are pairwise not T-equivalent remain not $T(a)$-equivalent when viewed as $\mathcal{L}(a)$-formulas. So there is some $m \in \omega$ and some non-isolated type $q(y) \in S_m(T(a))$, which is omitted in $M_{1/2}(a)$.

Let $M_{1/2}$ be the reduct of $M_{1/2}(a)$ to $\mathcal{L}$. Then $M_{1/2}$ realizes $p(x)$ (by $a' = a^{M_{1/2}}$), so it is not isomorphic to M_0. It is clear that $M_{1/2}(a) \not\cong M_1(a)$, but we need to show $M_{1/2} \not\cong M_1$. Note that $M_{1/2}$ omits $q(y)$, viewed as a type over the finitely many parameters a'. Suppose for contradiction that $f\colon M_{1/2} \to M_1$ is an isomorphism. Then M_1 omits f_*q, which is a complete type over $f(a') \in M_1^n$, contradicting the saturation of M_1. □

Returning to theories with infinitely many countable models, we have seen examples where $I(T, \aleph_0) = \aleph_0$ and $I(T, \aleph_0) = 2^{\aleph_0}$. Those interested in set theory will wonder about the possibility of cardinals between $\aleph_0$ and $2^{\aleph_0}$.

Conjecture 7.37 (Vaught)**.** There is no theory T such that

$$\aleph_0 < I(T, \aleph_0) < 2^{\aleph_0}.$$

Vaught's conjecture is one of the oldest open problems in model theory. Note that if we assume the continuum hypothesis, it is trivial, since there are no cardinals between $\aleph_0$ and $2^{\aleph_0} = \aleph_1$. The question is whether it is possible to prove Vaught's conjecture from the axioms of ZFC.

Using tools from descriptive set theory, Morley came close to settling the conjecture.

Theorem 7.38 (Morley)**.** *There is no complete theory T such that*

$$\aleph_1 < I(T, \aleph_0) < 2^{\aleph_0}.$$

In light of the theorems of Vaught and Morley, the possibilities for $I(T, \aleph_0)$ when T is complete are: $1, 3, 4, 5, \ldots, \aleph_0, \aleph_1, 2^{\aleph_0}$. And Vaught's conjecture concerns the open case of $\aleph_1$.

In the discussion above, we have only considered $I(T, \kappa)$ for $\kappa = \aleph_0$. The case of uncountable κ is a whole other story, beginning with another theorem of Morley's (the categoricity theorem) and extended by Shelah in his landmark book *Classification Theory.*