Slides for a course on model theory

Benno van den Berg

Utrecht University, Fall 2012

Quick history of logic

- Aristotle (384-322 BC): idea of *formal logic*. Syllogisms.
- Chryssipus (mid 3rd century BC): propositional logic.
- Frege (1848-1924): quantifiers, first-order logic.
- Gödel (1906-1978): completeness theorem.

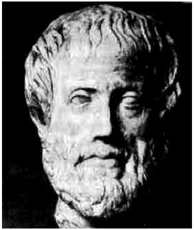
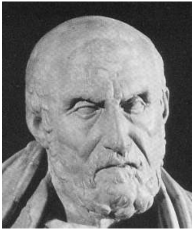

Tarski and Robinson

Founding father of model theory: Alfred Tarski (1901-1983). Created a school in Berkeley in the sixties.

Another important name is Abraham Robinson (1918-1974).




Stability theory

Morley's Theorem (1965): starting point for stability theory.

Shelah: classification theory.

More applied direction (geometric stability theory): Zil'ber and Hrushovski.

Applications

- 1968: Ax-Kochen-Ershov proof of Artin's conjecture.
- 1993: Hrushovski's proof of the Mordell-Lang conjecture for function fields.
- 2009: Pila's work on the Andre-Oort conjecture.

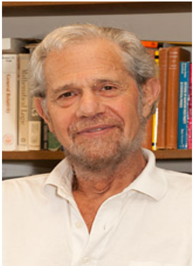


Literature

- Wilfrid Hodges, *A shorter model theory*. CUP 1997.
- David Marker, *Model theory: an introduction*. Springer 2002.
- Tent and Ziegler. *A course in model theory*. Lecture Notes in Logic, 2012.

Free internet sources:

- Achim Blumensath, *Logic, algebra and geometry*. http://www.mathematik.tu-darmstadt.de/∼blumensath/
- Jaap van Oosten, lecture notes for a course given in Spring 2000. http://www.staff.science.uu.nl/∼ooste110/syllabi/modelthmoeder.pdf
- C. Ward Henson, lecture notes for a course given in Spring 2010. http://www.math.uiuc.edu/∼henson/Math571/Math571Spring2010.pdf

We will not cover *finite model theory*. For that see

- Ebbinghaus and Flum, *Finite model theory*. Springer, 1995.

Language

A *language* or *signature* consists of:

1. constants.
2. function symbols.
3. relation symbols.

Once and for all, we fix a countably infinite set of variables. The terms are the smallest set such that:

1. all constants are terms.
2. all variables are terms.
3. if $t_1, \ldots, t_n$ are terms and f is an n-ary function symbol, then also $f(t_1, \ldots, t_n)$ is a term.

Terms which do not contain any variables are called *closed*.

Formulas and sentences

The *atomic formulas* are:

1. $s = t$, where s and t are terms.
2. $P(t_1, \ldots, t_n)$, where $t_1, \ldots, t_n$ are terms and P is a predicate symbol.

The set of *formulas* is the smallest set which:

1. contains the atomic formulas.
2. is closed under the propositional connectives $\wedge, \vee, \rightarrow, \neg$.
3. contains $\exists x\, \varphi$ and $\forall x\, \varphi$, if φ is a formula.

A formula which does not contain any quantifiers is called *quantifier-free*. A *sentence* is a formula which does not contain any free variables. A set of sentences is called a *theory*.

Convention: If we write $\varphi(x_1, \ldots, x_n)$, this is supposed to mean: φ is a formula and its free variables are contained in $\{x_1, \ldots, x_n\}$.

Models

A *structure* or *model* M in a language L consists of:

1. a set M (the *domain* or the *universe*).
2. interpretations $c^M \in M$ of all the constants in L,
3. interpretations $f^M : M^n \to M$ of all function symbols in L,
4. interpretations $R^M \subseteq M^n$ of all relation symbols in L.

The interpretation can then be extended to all terms in the language:

$$f(t_1, \ldots, t_n)^M = f^M(t_1^M, \ldots, f_n^M).$$

Tarski's truth definition

Let M be a model in a language L. Let L_M be the language obtained by adding fresh constants $\{c_m : m \in M\}$ to the language L, with c_m to be interpreted as m. We will seldom distinguish between c_m and m.

Validity or truth

If M is a model and φ is a sentence in the language L_M, then:

- $M \models s = t$ iff $s^M = t^M$;
- $M \models P(t_1, \dots, t_n)$ iff $(t_1, \dots, t_n) \in P^M$;
- $M \models \varphi \wedge \psi$ iff $M \models \varphi$ and $M \models \psi$;
- $M \models \varphi \vee \psi$ iff $M \models \varphi$ or $M \models \psi$;
- $M \models \varphi \rightarrow \psi$ iff $M \models \varphi$ implies $M \models \psi$;
- $M \models \neg\varphi$ iff not $M \models \varphi$;
- $M \models \exists x\, \varphi(x)$ iff there is an $m \in M$ such that $M \models \varphi(m)$;
- $M \models \forall x\, \varphi(x)$ iff for all $m \in M$ we have $M \models \varphi(m)$.

Semantic implication

Definition

If M is a model in a language L, then $\mathrm{Th}(M)$ is the collection L-sentences true in M. If N is another model in the language L, then we write $M \equiv N$ and call M and N *elementary equivalent*, whenever $\mathrm{Th}(M) = \mathrm{Th}(N)$.

Definition

Let Γ and Δ be theories. If $M \models \varphi$ for all $\varphi \in \Gamma$, then M is called a *model* of Γ. We will write $\Gamma \models \Delta$ if every model of Γ is a model of Δ as well. We write $\Gamma \models \varphi$ for $\Gamma \models \{\varphi\}$, et cetera.

Expansions and reducts

If $L \subseteq L'$ and M is an L'-structure, then we can obtain an L-structure N by taking the universe of M and forgetting the interpretations of the symbols which do not occur in L. In that case, M is an *expansion* of N and N is the *L-reduct* of M.

Lemma

If $L \subseteq L'$ and M is an L-structure and N is its L-reduct, then we have $N \models \varphi(m_1, \ldots, m_n)$ iff $M \models \varphi(m_1, \ldots, m_n)$ for all formulas $\varphi(x_1, \ldots, x_n)$ in the language L.

Homomorphisms

Let M and N be two L-structures. A *homomorphism* $h : M \to N$ is a function $h : M \to N$ such that:

1. $h(c^M) = c^N$ for all constants c in L;
2. $h(f^M(m_1, \ldots, m_n)) = f^N(h(m_1), \ldots, h(m_n))$ for all function symbols f in L and elements $m_1, \ldots, m_n \in M$;
3. $(m_1, \ldots, m_n) \in R^M$ implies $(h(m_1), \ldots, h(m_n)) \in R^N$.

A homomorphism which is bijective and whose inverse f^{-1} is also a homomorphism is called an *isomorphism*. If an isomorphism exists between structures M and N, then M and N are called *isomorphic*. An isomorphism from a structure to itself is called an *automorphism*.

Embeddings

A homomorphism $h : M \to N$ is an *embedding* if

1. h is injective;
2. $(h(m_1), \dots, h(m_n)) \in R^N$ implies $(m_1, \dots, m_n) \in R^M$.

Lemma

The following are equivalent for a homomorphism $h : M \to N$:

1. it is an embedding.
2. $M \models \varphi(m_1, \dots, m_n) \Leftrightarrow N \models \varphi(h(m_1), \dots, h(m_n))$ for all $m_1, \dots, m_n \in M$ and atomic formulas $\varphi(x_1, \dots, x_n)$.
3. $M \models \varphi(m_1, \dots, m_n) \Leftrightarrow N \models \varphi(h(m_1), \dots, h(m_n))$ for all $m_1, \dots, m_n \in M$ and quantifier-free formulas $\varphi(x_1, \dots, x_n)$.

If M and N are two models and the inclusion $M \subseteq N$ is an embedding, then M is a *substructure* of N and N is an *extension* of M.

Elementary embeddings

An embedding is called *elementary*, if

$$M \models \varphi(m_1, \ldots, m_n) \Leftrightarrow N \models \varphi(h(m_1), \ldots, h(m_n))$$

for all $m_1, \ldots, m_n \in M$ and *all* formulas $\varphi(x_1, \ldots, x_n)$.

Lemma

If h is an isomorphism, then h is an elementary embedding. If there is an elementary embedding $h : M \to N$, then $M \equiv N$.

Tarski-Vaught Test

If $h : M \to N$ is an embedding, then it is elementary iff for any formula $\varphi(y, x_1, \ldots, x_k)$ and $m_1, \ldots, m_k \in M$ and $n \in N$ such that $N \models \varphi(n, h(m_1), \ldots, h(m_k))$, there is an $m \in M$ such that $N \models \varphi(h(m), h(m_1), \ldots, h(m_k))$.

Recap on cardinal numbers

Two sets X and Y are *equinumerous* if there is a bijection from X to Y. Equinumerosity is an equivalence relation. For every set X there is an equinumerous set $|X|$ such that X and Y are equinumerous iff $|X| = |Y|$. A set of the form $|X|$ is called a *cardinal number* and $|X|$ is the *cardinality* of X. We will use small Greek letters $\kappa, \lambda \ldots$ for cardinal numbers.

We write $\kappa \leq \lambda$ if there is an injection from κ to λ. This gives the cardinals numbers the structure of a linear order. In fact, it is a well-order: every non-empty class of cardinal numbers has a least element.

Recap on cardinal numbers, continued

The smallest infinite cardinal number is $|\mathbb{N}|$, often written $\aleph_0$ or ω. Sets which have this cardinality are called *countably infinite*. Smaller sets are *finite* and bigger sets *uncountable*. A set which is either finite or countably infinite is called *countable*.

The cardinality of $2^{\mathbb{N}}$ is often called the continuum. The continuum hypothesis says it is smallest uncountable cardinal.

Recap on cardinal numbers, continued

Cardinal arithmetic is easy: define $\kappa + \lambda$ to be the cardinality of disjoint union of κ and λ and $\kappa \cdot \lambda$ to be the cardinality of the cartesian product of κ and λ. Then we have

$$\kappa + \lambda = \kappa \cdot \lambda = \max(\kappa, \lambda)$$

if at least one of κ, λ is infinite. Of course, cardinal exponentiation is hard!

If X is an infinite set, then X and the collection of finite subsets of X have the same cardinality.

Cardinality of model and language

Definition

The *cardinality* of a model is the cardinality of its underlying domain. The cardinality of a language L is the sums of the cardinalities of its sets of constants, function symbols and relation symbols.

Universal theories

Universal theory

A sentence is *universal* if it starts with a string of universal quantifiers followed by a quantifier-free formula. A theory is *universal* if it consists of universal sentences. A theory has a *universal axiomatisation* if it has the same class of models as a universal theory in the same language.

Examples of theories which have a universal axiomatisation:

- Groups
- Rings
- Commutative rings
- Vector spaces
- Directed and undirected graphs

Non-example:

- Fields

Exercises

Proposition

If T has a universal axiomatisation, then its class of models is closed under substructures.

Proof.

Exercise! (Challenge: Prove the converse!) □

Proposition

The theory of fields has no universal axiomatisation.

Proof.

Exercise! □

Skolem's Theorem

Theorem (Skolem)

Let L be a language. Then there is a language $L' \supseteq L$ with $|L'| \leq |L| + \aleph_0$ and a universal theory Sk_L in the language L' such that:

1. every L-formula is equivalent over Sk_L to a quantifier-free L'-formula.
2. every L-structure has an expansion to an L'-structure which is a model of Sk_L.

Proof.

For every quantifier-free formula $\varphi(x_1, \ldots, x_n, y)$ in the language L with at least one free variable we add to L' the n-ary function symbol f_φ and to Sk_L the universal sentence

$$\forall x_1, \ldots, x_n \forall y \left(\varphi(x_1, \ldots, x_n, y) \to \varphi(x_1, \ldots, x_n, f_\varphi(x_1, \ldots, x_n)) \right).$$

□

Skolem theories

Definition

An L-theory T is a *Skolem theory* or *has built-in Skolem functions* if for every formula $\varphi(x_1, \ldots, x_n, y)$ there is a function symbol f such that

$$T \models \forall x_1, \ldots, x_n \left(\exists y \, \varphi(x_1, \ldots, x_n, y) \to \varphi(x_1, \ldots, x_n, f(x_1, \ldots, x_n)) \right).$$

It is sufficient to require this for quantifier-free φ. (Exercise!)

Theorem

For every theory T in a language L there is a Skolem theory $T' \supseteq T$ in a language $L' \supseteq L$ with $|L'| \leq |L| + \aleph_0$ such that every model of T has an expansion to a model of T'.

Proof.

Write $L_0 = L$. Then let L_{n+1} be the language of Sk_{L_n} and put $L' = \bigcup L_n$ and $T' = T \cup \bigcup \mathrm{Sk}_{L_n}$. □

A theory T' as in the theorem is called a *skolemisation* of T.

Skolem hulls

Let M be a model of a Skolem theory T. Then for every subset $X \subseteq M$ the smallest subset of M containing X and closed under all the interpretations of the function symbols can be given the structure of a submodel of M. This is called the *Skolem hull* generated by X and denoted by $\langle X \rangle$.

Proposition

$\langle X \rangle$ is an elementary substructure of M.

Proof.

Exercise! (Hint: use Tarski-Vaught.) □

Downward Löwenheim-Skolem

Downward Löwenheim-Skolem

Suppose M is an L-structure and $X \subseteq M$. Then there is an elementary substructure N of M with $X \subseteq N$ and $|N| \leq |X| + |L| + \aleph_0$.

Proof.

Let T be the skolemisation of the empty theory in the language L and M' the expansion of M to a model of T. Then let N' be the Skolem hull generated by X. Then N' is an elementary substructure of M', and the reduct N of N' to the language L is an elementary substructure of M. □

Exercises

Proposition

A Skolem theory has a universal axiomatisation.

Proof.

Exercise! □

Proposition

A Skolem theory has quantifier-elimination.

Proof.

Exercise! □

Compactness Theorem

Definition

A theory T is *consistent* if every finite subset of T has a model.

Compactness Theorem

If a theory in a language L is consistent, then it has a model of cardinality $\leq |L| + \aleph_0$.

We will first prove this for universal theories.

Compactness theorem for universal theories

Compactness theorem for universal theories

If a universal theory in a language L is consistent, then it has a model of cardinality $\leq |L| + \aleph_0$.

Proof. Let T be a universal theory in a language L which is consistent. Without loss of generality, we may assume that L contains at least one constant: otherwise, simply add one to the language.

Let Δ the set of literals in the language L (a *literal* is an atomic sentence or its negation). Then the set

$$\{\Gamma \subseteq \Delta \; : \; T \cup \Gamma \text{ is consistent } \}$$

is partially ordered by inclusion. Moreover, every chain has an upper bound, so it contains a maximal element Γ_0 by Zorn's Lemma. For every atomic sentence we have either $\varphi \in \Gamma_0$ or $\neg\varphi \in \Gamma_0$.

Proof continued

We are now going to create a model M on the basis of the set Γ_0. Let $\mathcal{T}$ be the collection of terms in the language L. On $\mathcal{T}$ we can define a relation by:

$$s \sim t \Leftrightarrow s = t \in \Gamma_0.$$

This is an equivalence relation.

We can now define the interpretation of constants, function and relation symbols, as follows:

$$\begin{array}{rcl} c^M & = & [c], \\ f^M([t_1], \ldots, [t_n]) & = & [f(t_1, \ldots, t_n)], \\ R^M([t_1], \ldots, [t_n]) & \Leftrightarrow & R(t_1, \ldots, t_n) \in \Gamma_0. \end{array}$$

Check that this is well-defined! We have for every term t that $t^M = [t]$. Moreover, the set of literals true in M coincides precisely with Γ_0.

Proof finished

In order to finish the proof we need to show that M is a model of T. So consider a universal sentence $\forall x_1 \ldots \forall x_n \, \psi(x_1, \ldots, x_n)$ (ψ quantifier-free) that belongs to T. To show that it is valid in M we need to prove that for all terms $t_1, \ldots, t_n$ we have

$$M \models \psi([t_1], \ldots, [t_n]), \text{ or } M \models \psi(t_1, \ldots, t_n).$$

Let S be the collection of all sentences all whose terms and relation symbols also occur in $\psi(t_1, \ldots, t_n)$ and put $\Gamma_1 = \Gamma_0 \cap S$. Since there occur only finitely many terms and relation symbols in $\psi(t_1, \ldots, t_n)$, the set Γ_1 is finite.

Because the set $T \cup \Gamma_0$ is consistent, there is a model N of $\{\forall x_1 \ldots \forall x_n \, \psi(x_1, \ldots, x_n)\} \cup \Gamma_1$. We have $N \models \varphi$ iff $\varphi \in \Gamma_1$ for all literals φ in S and hence $N \models \varphi$ iff $M \models \varphi$ for all quantifier-free sentences φ in S. So since we have $N \models \psi(t_1, \ldots, t_n)$, we have $M \models \psi(t_1, \ldots, t_n)$ as well. $\square$

Reduction

Lemma

Let T be a consistent theory in a language L. Then there is a language $L' \supseteq L$ with $|L'| \leq |L| + \aleph_0$ and a consistent universal theory T' in the language L' such that

1. every L-structures modelling T has an expansion to an L'-structure modelling T', and
2. every L-reduct of a model of T' is a model of T.

Proof.

Let L' be the language of Sk_L. By Skolem's theorem every sentence $\varphi \in T$ is equivalent modulo Sk_L to a quantifier-free sentence φ' in the language L'. Then let $T' = Sk_L \cup \{\varphi' : \varphi \in T\}$. □

General case

Compactness Theorem

If a theory in a language L is consistent, then it has a model of cardinality $\leq |L| + \aleph_0$.

Proof.

If T is a theory in language L which is consistent, then there is a universal theory T' in a richer language L' which is also consistent and is such that every L-reduct of a model of T' is a model of T. By the compactness theorem for universal theories, T' has a model M'. So the reduct of M' to L is a model of T. □

Diagrams

Definition

A *literal* is an atomic sentence or the negation of an atomic sentence. If M is a model in a language L, then the collection of L_M-literals true in M is called the *diagram* of M and written $\mathrm{Diag}(M)$. The collection of all L_M-sentences true in M is called the *elementary diagram* of M and written $\mathrm{ElDiag}(M)$.

Lemma

The following amount to the same thing:

- A model N of $\mathrm{Diag}(M)$.
- An embedding $h : M \to N$.

As do the following:

- A model N of $\mathrm{ElDiag}(M)$.
- An elementary embedding $h : M \to N$.

Upward Löwenheim-Skolem

Upward Löwenheim-Skolem

Suppose M is an infinite L-structure and κ is a cardinal number with $\kappa \geq |M|, |L|$. Then there is an elementary embedding $i : M \to N$ with $|N| = \kappa$.

Proof.

Let Γ be the elementary diagram of M and Δ be the set of sentences $\{c_i \neq c_j : i \neq j \in \kappa\}$ where the c_i are κ-many fresh constants. By the Compactness Theorem, the theory $\Gamma \cup \Delta$ has a model A; we have $|A| \geq \kappa$. By the downwards version A has an elementary substructure N of cardinality κ. So, since N is a model of Γ, there is an elementary embedding $i : M \to N$. □

Characterisation universal theories

Theorem

T has a universal axiomatisation iff models of T are closed under substructures.

Proof.

Suppose T is a theory such that its models are closed under substructures. Let $T' = \{\varphi : T \models \varphi$ and φ is universal $\}$. Clearly, $T \models T'$. We need to prove the converse.

So suppose M is a model of T'. It sufffices to show that $T \cup \mathrm{Diag}(M)$ is consistent. Because once we do that, it will have a model N. But since N is a model of $\mathrm{Diag}(M)$, it will be an extension of M; and because N is a model of T and models of T are closed under substructures, M will be a model of T. □

Proof of claim

Claim

If $M \models T'$ where $T' = \{\varphi : T \models \varphi$ and φ is universal $\}$, then $T \cup \mathrm{Diag}(M)$ is consistent.

Proof.

Suppose not. Then, by the compactness theorem, there would be a finite set of literals $\psi_1, \ldots, \psi_n \in \mathrm{Diag}(M)$ which are inconsistent with T. Replace the constants from M in $\psi_1, \ldots, \psi_n$ by variables $x_1, \ldots, x_n$ and we obtain $\psi'_1, \ldots, \psi'_n$; because the constants from M do not appear in T, the theory T is already inconsistent with $\exists x_1, \ldots, x_n\, (\, \psi'_1 \wedge \ldots, \wedge \psi'_n\,)$. But then it follows that $T \models \neg\exists_1, \ldots, x_n\, (\, \psi'_1 \wedge \ldots \psi'_n\,)$ and $T \models \forall x_1, \ldots, x_n\, (\, \neg(\psi'_1 \wedge \ldots \psi'_n)\,)$, and hence $\forall x_1, \ldots, x_n\, (\, \neg(\psi'_1 \wedge \ldots \psi'_n)\,) \in T'$. But this contradicts the fact that M is a model of T'. □

Two exercises

Exercise

Prove: a theory has an existential axiomatisation iff its models are closed under extensions.

Exercise

For two L-structures A and B, we have $A \equiv B$ iff A and B have a common elementary extension.

Directed systems

See Chapters IV-VI in the lecture notes by Jaap van Oosten.

Definition

A partially ordered set $(K, \leq)$ is called *directed*, if K is non-empty and for any two elements $x, y \in K$ there is an element $z \in K$ such that $x \leq z$ and $y \leq z$.

Definition

A *directed system* of L-structures consists of a family $(M_k)_{k \in K}$ of L-structures indexed by K, together with homomorphisms $f_{kl} : M_k \to M_l$ for $k \leq l$. These homomorphisms should satisfy:

- f_{kk} is the identity homomorphism on M_k,
- if $k \leq l \leq m$, then $f_{km} = f_{lm} f_{kl}$.

If we have a directed system, then we can construct its *colimit*.

The colimit

First, we take the disjoint union of all the universes:

$$\sum_{k \in K} M_k = \{(k, a) \,:\, k \in K, a \in M_k\},$$

and then we define an equivalence relation on it:

$$(k, a) \sim (l, b) :\Leftrightarrow (\exists m \geq k, l)\, f_{km}(a) = f_{lm}(b).$$

Let M be the set of equivalence classes and denote the equivalence class of (k, a) by $[k, a]$.

The colimit, continued

M has an L-structure: we put

$$f^M([k_1, a_1], \ldots, [k_n, a_n]) \quad = \quad [k, f^{M_k}(f_{k_1 k}(a_1), \ldots, f_{k_n k}(a_n)],$$

where k is an element $\geq k_1, \ldots, k_n$. (Check that this makes sense!)

And we put

$$R^M([k_1, a_1], \ldots, [k_n, a_n])$$

iff there is a $k \geq k_1, \ldots, k_n$ such that

$$(f_{k_1 k}(a_1), \ldots, f_{k_n k}(a_n)) \in R^{M_k}.$$

In addition, we have maps $f_k : M_k \to M$ sending a to $[k, a]$.

Omnibus theorem

The following theorem collects the most important facts about colimits of filtered systems. Especially useful is part 5.

Theorem

1. All f_k are homomorphisms.
2. If $k \leq l$, then $f_l f_{kl} = f_k$.
3. If N is another L-structure for which there are homomorphisms $g_k : M_k \to N$ such that $g_l f_{kl} = g_k$ whenever $k \leq l$, then there is a unique homomorphisms $g : M \to N$ such that $g f_k = g_k$ for all $k \in K$ ("universal property").
4. If all maps f_{kl} are embeddings, then so are all f_k.
5. If all maps f_{kl} are elementary embeddings, then so are all f_k ("elementary system lemma").

Proof.

Exercise! □

Next goal

Our next big goal will be to prove:

Robinson's Consistency Theorem

Let L_1 and L_2 be two languages and $L = L_1 \cap L_2$. Suppose T_1 is an L_1-theory, T_2 an L_2-theory and both extend a complete L-theory T. If both T_1 and T_2 are consistent, then so is $T_1 \cup T_2$.

We first treat the special case where $L_1 \subseteq L_2$.

First lemma

Lemma

Let $L \subseteq L'$, A an L-structure and B an L'-structure. Suppose moreover $A \equiv B \upharpoonright L$. Then there is an L'-structure C and a diagram of elementary embeddings (f in L and f' in L')

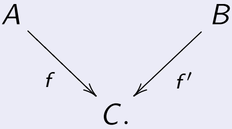

Proof. Consider $T = \mathrm{ElDiag}(A) \cup \mathrm{ElDiag}(B)$ (making sure we use different constants for the elements from A and B!). We need to show T has a model; so suppose T is inconsistent. Then, by Compactness, a finite subset of T has no model; taking conjunctions, we have sentences $\varphi(a_1, \ldots, a_n) \in \mathrm{ElDiag}(A)$ and $\psi(b_1, \ldots, b_m) \in \mathrm{ElDiag}(B)$ that are contradictory. But as the a_j do not occur in L_B, we must have that $B \models \neg\exists x_1, \ldots, x_n \, \varphi(x_1, \ldots, x_n)$. This contradicts $A \equiv B \upharpoonright L$. □

Second lemma

Lemma

Let $L \subseteq L'$ be languages, suppose A and B are L-structures and C is an L'-structure. Any pair of L-elementary embeddings $f : A \to B$ and $g : A \to C$ fit into a commuting square

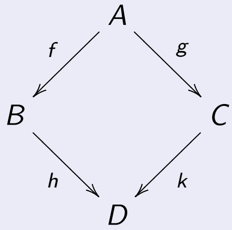

where D is an L'-structure, h is an L-elementary embedding and k is an L'-elementary embedding.

Proof.

Without loss of generality we may assume that L contains constants for all elements of A. Then simply apply the first lemma. □

Robinson's consistency theorem

Theorem

Let L_1 and L_2 be two languages and $L = L_1 \cap L_2$. Suppose T_1 is an L_1-theory, T_2 an L_2-theory and both extend a complete L-theory T. If both T_1 and T_2 are consistent, then so is $T_1 \cup T_2$.

Proof. Let A_0 be a model of T_1 and B_0 be a model of T_2. Since T is complete, their reducts to L are elementary equivalent, so, by the first lemma, there is a diagram

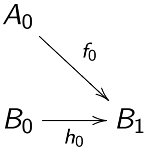

with h_0 an L_2-elementary embedding and f_0 an L-elementary embedding. Now by applying the second lemma to f_0 and the identity on A, we obtain

...

Robinson's consistency theorem, proof finished

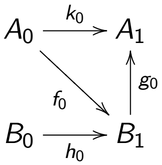

where g_0 is L-elementary and k_0 is L_1-elementary. Continuing in this way we obtain a diagram

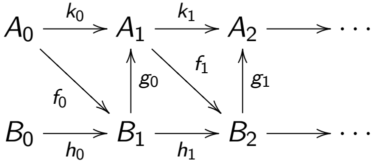

where the k_i are L_1-elementary, the f_i and g_i are L-elementary and the h_i are L_2-elementary. The colimit C of this directed system is both the colimit of the A_i and of the B_i. So A_0 and B_0 embed elementarily into C by the elementary systems lemma; hence C is a model of both T_1 and T_2, as desired. □

Amalgamation Theorem

Amalgamation Theorem

Let L_1, L_2 be languages and $L = L_1 \cap L_2$, and suppose A, B and C are structures in the languages L, L_1 and L_2, respectively. Any pair of L-elementary embeddings $f : A \to B$ and $g : A \to C$ fit into a commuting square

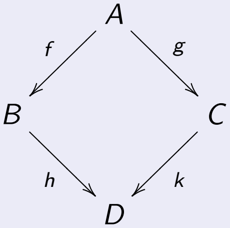

where D is an $L_1 \cup L_2$-structure, h is an L_1-elementary embedding and k is an L_2-elementary embedding.

Proof.

Immediate consequence of Robinson's Consistency Theorem. (Why?) □

Craig Interpolation

Craig Interpolation Theorem

Let φ and ψ be sentences in some language such that $\varphi \models \psi$. Then there is a sentence θ such that

1. $\varphi \models \theta$ and $\theta \models \psi$;
2. every predicate, function or constant symbol that occurs in θ occurs also in both φ and ψ.

Proof.

Let L be the common language of φ and ψ. We will show that $T_0 \models \psi$ where $T_0 = \{\sigma \in L : \varphi \models \sigma\}$. This is sufficient: for then there are $\theta_1, \ldots, \theta_n \in T_0$ such that $\theta_1, \ldots, \theta_n \models \psi$ by Compactness. So $\theta := \theta_1 \wedge \ldots \wedge \theta_n$ is the interpolant. □

Craig Interpolation, continued

Lemma

Let L be the common language of φ and ψ. If $\varphi \models \psi$, then $T_0 \models \psi$ where $T_0 = \{\sigma \in L : \varphi \models \sigma\}$.

Proof.

Suppose not. Then $T_0 \cup \{\neg\psi\}$ has a model A. Write $T = \mathrm{Th}_L(A)$. We now have $T_0 \subseteq T$ and:

1. T is a complete L-theory.
2. $T \cup \{\neg\psi\}$ is consistent (because A is a model).
3. $T \cup \{\varphi\}$ is consistent.

(Proof of 3: Suppose not. Then, by Compactness, there would a sentence $\sigma \in T$ such that $\varphi \models \neg\sigma$. But then $\neg\sigma \in T_0 \subseteq T$. Contradiction!)

Now we can apply Robinson's Consistency Theorem to deduce that $T \cup \{\neg\psi, \varphi\}$ is consistent. But that contradicts $\varphi \models \psi$. □

Beth Definability Theorem

Definition

Let L be a language a P be a predicate symbol not in L, and let T be an $L \cup \{P\}$-theory. T *defines* P *implicitly* if any L-structure M has at most one expansion to an $L \cup \{P\}$-structure which models T. There is another way of saying this: let T' be the theory T with all occurrences of P replaced by P'. Then T *defines* P *implicitly* iff

$$T \cup T' \models \forall x_1, \dots x_n \left(P(x_1, \dots, x_n) \leftrightarrow P'(x_1, \dots, x_n) \right).$$

T *defines* P *explicitly*, if there is an L-formula $\varphi(x_1, \dots, x_n)$ such that

$$T \models \forall x_1, \dots, x_n \left(P(x_1, \dots, x_n) \leftrightarrow \varphi(x_1, \dots, x_n) \right).$$

Beth Definability Theorem

T defines P implicitly if and only if T defines P explicitly.

(Right-to-left direction is obvious.)

Beth Definability Theorem, proof

Proof. Suppose T defines P implicitly. Add new constants $c_1, \ldots, c_n$ to the language. Then we have $T \cup T' \models P(c_1, \ldots, c_n) \to P'(c_1, \ldots, c_n)$. By Compactness and taking conjunctions we can find an $L \cup \{P\}$-formula ψ such that $T \models \psi$ and

$$\psi \wedge \psi' \models P(c_1, \ldots, c_n) \to P'(c_1, \ldots, c_n)$$

(where ψ' is ψ with all occurrences of P replaced by P'). Taking all the Ps to one side and the P's to another, we get

$$\psi \wedge P(c_1, \ldots, c_n) \models \psi' \to P'(c_1, \ldots, c_n)$$

So there is a Craig Interpolant θ such that

$$\psi \wedge P(c_1, \ldots, c_n) \models \theta \text{ and } \theta \models \psi' \wedge P'(c_1, \ldots, c_n)$$

By symmetry also

$$\psi' \wedge P'(c_1, \ldots, c_n) \models \theta \text{ and } \theta \models \psi \wedge P(c_1, \ldots, c_n)$$

So $\theta = \theta(c_1, \ldots, c_n)$ is, modulo T, equivalent to $P(c_1, \ldots, c_n)$ and $\theta(x_1, \ldots, x_n)$ defines P explicitly. $\Box$

Chang-Łoś-Suszko Theorem

Definition

A Π_2-*sentence* is a sentence which consists first of a sequence of universal quantifiers, then a sequence of existential quantifers and then a quantifier-free formula.

Definition

A theory T is *preserved by directed unions* if, for any directed system consisting of models of T and embeddings between them, also the colimit is a model T.

Chang-Łoś-Suszko Theorem

A theory is preserved under directed unions if and only if T can be axiomatised by Π_2-sentences.

Proof.

The easy direction is: Π_2-sentences are preserved by directed unions. We do the other direction. □

Chang-Łoś-Suszko Theorem, proof

Proof. Suppose T is preserved by direction unions. Again, let

$$T_0 = \{\varphi : \varphi \text{ is } \Pi_2 \text{ and } T \models \varphi\},$$

and let B be a model of T_0. We will construct a directed chain of embeddings

$$B = B_0 \to A_0 \to B_1 \to A_1 \to B_2 \to A_2 \ldots$$

such that:

1. Each A_n is a model of T.
2. The composed embeddings $B_n \to B_{n+1}$ are elementary.
3. Every universal sentence in the language L_{B_n} true in B_n is also true in A_n (when regarding A_n is an L_{B_n}-structure via the embedding $B_n \to A_n$).

This will suffice, because when we take the colimit of the chain, then it is:

- the colimit of the A_n, and hence a model of T, by assumption on T.
- the colimit of the B_n, and hence elementary equivalent to each B_n.

So B is a model of T, as desired.

Chang-Łoś-Suszko Theorem, proof continued

Construction of A_n: We need A_n to be a model of T and every universal sentence in the language L_{B_n} true in B_n to be true in A_n as well. So let

$$T' = T \cup \{\varphi \in L_{B_n} : \varphi \text{ universal and } B_n \models \varphi\};$$

to show that T' is consistent. Suppose not. Then there is a universal sentence $\forall x_1, \dots x_n \, \varphi(x_1, \dots, x_n, b_1, \dots, b_k)$ with $b_i \in B_n$ that is inconsistent with T. So

$$T \models \exists x_1, \dots, x_n \neg\varphi(x_1, \dots, x_n, b_1, \dots, b_k)$$

and

$$T \models \forall y_1, \dots, y_k \, \exists x_1, \dots, x_n \neg\varphi(x_1, \dots, x_n, y_1, \dots, y_k)$$

because the b_i do not occur in T. But this contradicts the fact that B_n is a model of T_0.

Chang-Łoś-Suszko Theorem, proof finished

Construction of B_{n+1}: We need $A_n \to B_{n+1}$ to be an embedding and $B_n \to B_{n+1}$ to be elementary. So let

$$T' = \mathrm{Diag}(A_n) \cup \mathrm{ElDiag}(B_n)$$

(identifying the element of B_n with their image along the embedding $B_n \to A_n$); to show that T' is consistent. Suppose not. Then there is a quantifier-free sentence

$$\varphi(b_1, \ldots, b_n, a_1, \ldots, a_k)$$

with $b_i \in B_n$ and $a_i \in A_n \setminus B_n$ which is true in A_n, but is inconsistent with $\mathrm{ElDiag}(B_n)$. Since the a_i do not occur in B_n, we must have

$$B_n \models \forall x_1, \ldots, x_k \neg\varphi(b_1, \ldots, b_n, x_1, \ldots, x_k).$$

This contradicts the fact that all universal L_{B_n}-sentences true in B_n are also true in A_n. □

Types

Fix $n \in \mathbb{N}$ and let $x_1, \dots, x_n$ be a fixed sequence of distinct variables.

Definition

- A *partial n-type in L* is a collection of formulas $\varphi(x_1, \dots, x_n)$ in L.
- If A is an L-structure and $a_1, \dots, a_n \in A$, then the *type of* $(a_1, \dots, a_n)$ *in* A is the set of L-formulas

  $$\{\varphi(x_1, \dots, x_n) : A \models \varphi(a_1, \dots, a_n)\};$$

  we denote this set by $\mathrm{tp}_A(a_1, \dots, a_n)$ or simply by $\mathrm{tp}(a_1, \dots, a_n)$ if A is understood.
- A *n-type in L* is a set of formulas of the form $\mathrm{tp}_A(a_1, \dots, a_n)$ for some L-structure A and some $a_1, \dots, a_n \in A$.

Realizing and omitting types

Definition

- If $\Gamma(x_1, \ldots, x_n)$ is a partial n-type in L, we say $(a_1, \ldots, a_n)$ *realizes* Γ *in* A if every formula in Γ is true of $a_1, \ldots, a_n$ in A.
- If $\Gamma(x_1, \ldots, x_n)$ is a partial n-type in L and A is an L-structure, we say that Γ is *realized or satisfied in* A if there is some n-tuple in A that realizes Γ in A. If no such n-tuple exists, then we say that A *omits* Γ.
- If $\Gamma(x_1, \ldots, x_n)$ is a partial n-type in L and A is an L-structure, we say that Γ is *finitely satisfiable in* A if any finite subset of Γ is realized in A.

Exercises

Exercise

Show that a partial n-type is an n-type iff it is finitely satisfiable and contains $\varphi(x_1, \ldots, x_n)$ or $\neg\varphi(x_1, \ldots, x_n)$ for every L-formula φ whose free variables are among the fixed variables $x_1, \ldots, x_n$.

Exercise

Show that a partial n-type can be extended to an n-type iff it is satisfiable.

Exercise

Suppose $A \equiv B$. If $\Gamma(x_1, \ldots, x_n)$ is finitely satisfiable in A, then it is also finitely satisfiable in B.

Logic topology

Definition

Let T be a theory in L and let $\Gamma = \Gamma(x_1, \ldots, x_n)$ be a partial n-type in L.

- Γ *is consistent with* T if $T \cup \Gamma$ has a model.
- The set of all n-types consistent with T is denoted by $S_n(T)$. These are exactly the n-types in L that contain T.

The set $S_n(T)$ can be given the structure of a topological space, where the basic open sets are given by

$$[\varphi(x_1, \ldots, x_n)] = \{\Gamma(x_1, \ldots, x_n) \in S_n(T) : \varphi \in \Gamma\}.$$

This is called the *logic topology*.

Type spaces

Theorem

The space $S_n(T)$ with the logic topology is a totally disconnected, compact Hausdorff space. Its closed sets are the sets of the form

$$\{\Gamma \in S_n(T) : \Gamma' \subseteq \Gamma\}$$

where Γ' is a partial n-type. In fact, two partial n-types are equivalent over T iff they determine the same closed set. Furthermore, the clopen sets in the type space are precisely the ones of the form $[\varphi(x_1, \ldots, x_n)]$.

κ-saturated models

Let A be an L-structure and X a subset of A. We write L_X for the language L extended with constants for all elements of X and $(A, a)_{a \in X}$ for the L_X-expansion of A where we interpret the constant $a \in X$ as itself.

Definition

Let A be an L-structure and let κ be an infinite cardinal. We say that A is *κ-saturated* if the following condition holds: if X is any subset of A having cardinality $< \kappa$ and $\Gamma(x)$ is any 1-type in L_X that is finitely satisfiable in $(A, a)_{a \in X}$, then $\Gamma(x)$ is itself satisfied in $(A, a)_{a \in X}$.

Remark

1. If A is infinite and κ-saturated, then A has cardinality at least κ.
2. If A is finite, then A is κ-saturated for every κ.
3. If A is κ-saturated and X is a subset of A having cardinality $< \kappa$, then $(A, a)_{a \in X}$ is also κ-saturated.

Property of κ-saturated models

Theorem

Suppose κ is an infinite cardinal, A is κ-saturated and $X \subseteq A$ is a subset of cardinality $< \kappa$. Suppose $\Gamma(y_i : i \in I)$ is a collection of L_X-formulas with $|I| \leq \kappa$. If Γ is finitely satisfiable in $(A, a)_{a \in X}$, then Γ is satisfiable in $(A, a)_{a \in X}$.

Proof.

Without loss of generality we may assume that $I = \kappa$ and Γ is complete: contains either φ or $\neg\varphi$ for every L_X-formula φ with free variables among $\{y_i : i \in \kappa\}$.

Write $\Gamma_{\leq j}$ for the collection of those elements of Γ that only contain variables y_i with $i \leq j$. By induction on j we will find an element a_j such that $(a_i)_{i \leq j}$ realizes $\Gamma_{\leq j}$. Consider Γ' which is $\Gamma_{\leq j}$ with all y_i replaced by a_i for $i < j$. This is a 1-type which is finitely satisfiable in $(A, a)_{a \in X \cup \{a_i : i < j\}}$ (check!). Since $(A, a)_{a \in X \cup \{a_i : i < j\}}$ is κ-saturated, we find a suitable a_j. □

Other notions of richness

Definition

Let A and B be L-structures and $X \subseteq A$. A map $f : X \to B$ will be called an *elementary map* if

$$A \models \varphi(a_1, \ldots, a_n) \Leftrightarrow B \models \varphi(f(a_1), \ldots, f(a_n))$$

for all L-formulas φ and $a_1, \ldots, a_n \in X$.

Definition

A structure M is

- κ*-universal* if every structure of cardinality $< \kappa$ which is elementarily equivalent to M can be elementarily embedded into M.
- κ*-homogeneous* if for every subset A of M of cardinality smaller than κ and for every $b \in M$, every elementary map $A \to M$ can be extended to an elementary map $A \cup \{b\} \to M$.

More properties of κ-saturated models

Theorem

Let M be an L-structure and $\kappa \geq |L|$ be infinite. If M is κ-saturated, then M is κ^+-universal and κ-homogeneous.

Proof.

Let M be κ-structure. First suppose A is a structure with $A \equiv M$ and $|A| \leq \kappa$. Consider $\Gamma = \mathrm{ElDiag}(A)$. Since $A \equiv M$, the set Γ is finitely satisfiable in M. By the theorem two slides ago, Γ is satisfiable in M, so A embeds elementarily in M.

Now let A be a subset of M with $|A| < \kappa$, $b \in M$ and $f : A \to M$ be elementary. Consider $\Gamma = \mathrm{tp}_{(M,a)_{a \in A}}(b)$. Since $(M, a)_{a \in A} \equiv (M, f(a))_{a \in A}$, the type $\Gamma(x)$ is finitely satisfiable in $(M, f(a))_{a \in M}$. Hence it is satisfied in M by some $c \in M$. Extend f by $f(b) = c$. $\square$

Exercise

In fact we have:

Theorem

Let M be an L-structure and $\kappa \geq |L|$ be infinite. Then the following are equivalent:

(1) M is κ-saturated.

(2) M is κ^+-universal and κ-homogeneous.

If $\kappa > |L| + \aleph_0$, this is also equivalent to:

(3) M is κ-universal and κ-homogeneous.

Proof.

Exercise! (Please try!) □

Theorem on saturated models

Theorem

Let $\kappa \geq |L|$ be infinite. Any two κ-saturated models of cardinality κ that are elementarily equivalent are isomorphic.

Proof.

By a back-and-forth argument. Let A, B be two elementarily equivalent saturated models of cardinality κ. By induction on κ we construct an increasing sequence of elementary maps $f_\alpha : X_\alpha \to B$ with $\bigcup_\alpha X_\alpha = A$ and $\bigcup_\alpha f(X_\alpha) = B$. Then $f = \bigcup_\alpha f_\alpha$ will be our desired isomorphism.

We start with $f_0 = \emptyset$ and at limit stages we simply take the union. At successor stages we alternate: at odd stages α we take a fresh element $a \in A$ and extend the map so that $a \in X_\alpha$; at even stages we take a fresh element $b \in B$ and extend the map so that $b \in f(X_\alpha)$. □

Strong homogeneity

Definition

A model M is *strongly κ-homogeneous* if for every subset A of M of cardinality strictly less than κ, every elementary map $A \to M$ can be extended to an automorphism of M.

Corollary

Let $\kappa \geq |L|$ be infinite. A model of cardinality κ that is κ-saturated is strongly κ-homogeneous.

Proof.

Let $f : A \to M$ be an elementary map and $|A| < \kappa$. Then $(M, a)_{a \in A}$ and $(M, f(a))_{a \in A}$ are elementary equivalent. Since both are κ-saturated, they must be isomorphic by the previous result. This isomorphism is the desired automorphism extending f. □

Exercises

Let $\kappa \geq |L|$ be infinite.

Exercise

Show that a strongly κ-homogeneous model is κ-homogeneous.

Exercise

Any κ-homogeneous model of cardinality κ is strongly homogeneous.

But do they exist?

So κ-saturated models are very nice. But we haven't answered a basic question: do they even exist? They do. In fact we have:

Theorem

For every infinite cardinal number κ, every structure has a κ-saturated elementary extension.

But to prove this we need a bit more set theory.

Cofinality

Recall that:

- An ordinal is a set consisting of all smaller ordinals.
- Ordinals can be of two sorts: they are either successor ordinals or limit ordinals. (Depending on whether they have a immediate predecessor.)
- A cardinal κ is ordinal which is the smallest among those having the same cardinality as κ. An infinite cardinal is always a limit ordinal.

Definition

Let α be a limit ordinal. A set $X \subseteq \alpha$ is called *bounded* if there is a $\beta \in \alpha$ such that $x \leq \beta$ for all $x \in X$; otherwise it is *unbounded* or *cofinal*. The cardinality of the smallest unbounded set is called the *cofinality* of α and written $\mathrm{cf}(\alpha)$.

Note: $\omega \leq \mathrm{cf}(\alpha) \leq \alpha$ and $\mathrm{cf}(\alpha)$ is a cardinal.

Cofinal map

Definition

A map $f : \alpha \to \beta$ is *cofinal*, if it is increasing and its image is unbounded.

Lemma

1. There is a cofinal map $\mathrm{cf}(\alpha) \to \alpha$.
2. If $f : \alpha \to \beta$ is cofinal, then $\mathrm{cf}(\alpha) = \mathrm{cf}(\beta)$.
3. $\mathrm{cf}(\mathrm{cf}(\alpha)) = \mathrm{cf}(\alpha)$.

Definition

A cardinal number κ for which $\mathrm{cf}(\kappa) = \kappa$ is called *regular*. Otherwise it is called *singular*.

Note: $\mathrm{cf}(\alpha)$ is always regular.

Regular cardinals

Theorem

Let κ be a cardinal. Suppose λ is the least cardinal for which there is a family of sets $\{X_i : i \in \lambda\}$ such that $|\sum_{i\in\lambda} X_i| = \kappa$ and $|X_i| < \kappa$. Then $\lambda = \mathrm{cf}(\kappa)$.

Theorem

Infinite successor cardinals are always regular.

Proof.

Immediate from the previous theorem and the fact that $\kappa \cdot \kappa = \kappa$ for infinite cardinals κ. □

Goal

Recall our goal was to prove:

Theorem

For every infinite cardinal number κ, every structure has a κ-saturated elementary extension.

We first prove a lemma.

A lemma

Lemma

Let A be an L-structure. There exists an elementary extension B of A such that for every subset $X \subseteq A$, every 1-type in L_X which is finitely satisfied in $(A, a)_{a \in X}$ is realized in $(B, a)_{a \in X}$.

Proof.

Let $(\Gamma_i(x_i))_{i \in I}$ be the collection of all such 1-types and b_i be new constants. Then every finite subset of

$$\Gamma := \bigcup_{i \in I} \Gamma_i(b_i)$$

is satisfied in $(A, a)_{a \in A}$, so it has a model B. Since Γ contains $\mathrm{ElDiag}(A)$, the model A embeds into B. □

Existence of rich models

Theorem

For every infinite cardinal number κ, every structure has a κ-saturated elementary extension.

Proof.

Let A be an L-structure. We will build an elementary chain of L-structures $(A_i : i \in \kappa^+)$. We set $A_0 = A$, at successor stages we apply the previous lemma and at limit stages we take the colimit. Now let B be the colimit of the entire chain. We claim B is κ^+-saturated (which is more than we need).

So let $X \subseteq B$ be a subset of cardinality $< \kappa^+$ and $\Gamma(x)$ be a 1-type in L_X that is finitely satisfied in $(A, a)_{a \in X}$. Since κ^+ is regular, there is an $i \in \kappa^+$ such that $X \subseteq A_i$. And since A embeds elementarily into A_i, the type $\Gamma(x)$ is also finitely satisfied in $(A_i, a)_{a \in X}$. So it is realized in A_{i+1}, and therefore also in B, because A_{i+1} embeds elementarily into B. □

Even richer models

Now that we have this we can be even more ambitious:

Theorem

For every infinite cardinal number κ, every structure has a κ-saturated elementary extension all whose reducts are strongly κ-homogeneous.

We need a lemma:

Lemma

Suppose A is κ-saturated and B is an elementary substructure of A satisfying $|B| < \kappa$. Then any elementary map f between subsets of B can be extended to an elementary embedding of B into A.

Proof.

If $f : S \to B$ is the elementary mapping, then $(B, b)_{b \in S} \equiv (A, f(b))_{b \in S}$. Since $|S| < \kappa$, also $(A, f(b))_{b \in S}$ is κ-saturated und hence κ^+-universal. So $(B, b)_{b \in S}$ embeds elementarily into $(A, f(b))_{b \in S}$: so we have an elementary embedding of B into A extending f. □

Existence of very rich models

Theorem

For every infinite cardinal number κ, every structure has a κ-saturated elementary extension all whose reducts are strongly κ-homogeneous.

Proof.

Let A be an L-structure. Again, we will build an elementary chain of L-structures $(M_\alpha : \alpha \in \kappa^+)$. We set $M_0 = A$, at successor stages $\alpha + 1$ we take an $|M_\alpha|^+$-saturated elementary extension of M_α and at limit stages we take the colimit. Now let M be the colimit of the entire chain. We claim M is as desired.

Any subset of S of M that has cardinality $\leq \kappa$, must be a subset of some M_α (using again that κ^+ is regular). So M is κ^+-saturated. It remains to show that every reduct of M is strongly κ-homogeneous. □

Existence of very rich models, proof finished

Proof.

Let f be any mapping between subsets of M that is elementary, with domain and range having cardinality $< \kappa$. Again, domain and range will belong to some M_α. Without loss of generality we may assume that α is a limit ordinal. We extend f to a map $f_\alpha : M_\alpha \to M_{\alpha+1}$ using the lemma.

We will build maps f_β for all $\alpha \leq \beta < \kappa^+$ in such a way that f_β is an elementary embedding of M_β in $M_{\beta+1}$ and $f_{\beta+1}$ extends f_β^{-1}. It follows that $f_{\beta+2}$ extends f_β and that the union h over all f_β with β even is an automorphism of M.

The construction is: At limit stages we take unions over all previous even stages. And at successor stages we apply the lemma.

This argument works equally well for reducts of M. □

Definability

Definition

Let A be an L-structure and $R \subseteq A^n$ be a relation. The relation R is called *definable*, if there a formula $\varphi(x_1, \ldots, x_n)$ such that

$$R = \{(a_1, \ldots, a_n) \in A^n : A \models \varphi(a_1, \ldots, a_n)\}.$$

A homomorphism $f : A \to A$ *leaves* R *setwise invariant* if $\{(f(a_1), \ldots, f(a_n) : (a_1, \ldots, a_n) \in R\} = R$.

Proposition

Every elementary embedding from A to itself leaves all definable relations setwise invariant.

Definability results

Theorem

Let L be a language and P a predicate not in L. Suppose (A, R) is an ω-saturated $L \cup \{P\}$-structure and that A is strongly ω-homogeneous. Then the following are equivalent:

(1) R is definable in A.

(2) every automorphism of A leaves R setwise invariant.

Proof.

(1) $\Rightarrow$ (2) always holds, because automorphisms are elementary embeddings.

(2) $\Rightarrow$ (1): Suppose R is not definable. By the next lemma there are tuples a and b having the same type such that $R(a)$ is true and $R(b)$ is false. But then there is an automorphism of A that sends a to b by strong homogeneity. So R is not setwise invariant under automorphisms of A. □

A lemma

Lemma

Suppose A is a structure and R is not definable in A. If (A, R) is ω-saturated, then there are tuples a and b having the same n-type in A such that $R(a)$ is true and $R(b)$ is false.

Proof.

First consider the type
$\Sigma(x) = \{\varphi(x) \in L : (A, R) \models \forall x\,(\neg P(x) \to \varphi(x)\} \cup \{P(x)\}$. This type is finitely satisfiable in (A, R): for if not, then there would be a formula $\varphi(x)$ such that $(A, R) \models \neg P(x) \to \varphi(x)$ and $(A, R) \models \neg(\varphi(x) \wedge P(x))$. But then $\neg\varphi(x)$ would define R. By ω-saturation, there is an element a realizing $\Sigma(x)$. Now consider the type $\Gamma(x) = \mathrm{tp}_A(a) \cup \{\neg P(x)\}$. This type is also finitely satisfiable in (A, R): for if not, then there would be a formula $\varphi(x) \in L$ such that $(A, R) \models \varphi(a)$ and $(A, R) \models \neg(\varphi(x) \wedge \neg P(x))$. This is impossible by construction of a. By ω-saturation there is an element b realizing $\Gamma(x)$. So we have that a and b have the same type in A, while $R(a)$ is true and $R(b)$ is false. □

Svenonius' Theorem

Svenonius' Theorem

Let A be an L-structure and R be a relation on A. Then the following are equivalent:

(1) R is definable in A.

(2) every automorphism of an elementary extension (B, S) of (A, R) leaves S setwise invariant.

Proof.

(1) $\Rightarrow$ (2): If R is definable in A, then S is definable in B by the same formula; so it will be left setwise invariant by any automorphism.

(2) $\Rightarrow$ (1): Let (B, S) be an ω-saturated and strongly ω-homogeneous extension of (A, R). S will be definable in (B, S) by the previous theorem; but then R in A will be definable by the same formula. □

Omitting types theorem

Definition

Let T be an L-theory and $\Sigma(x)$ be a partial type. Then $\Sigma(x)$ is *isolated in* T if there is a formula $\varphi(x)$ such that $\exists x\, \varphi(x)$ is consistent with T and

$$T \models \varphi(x) \to \sigma(x)$$

for all $\sigma(x) \in \Sigma(x)$.

Exercise

A type is isolated iff it is an isolated point in the type space $S_1(T)$.

Omitting types theorem

Let T be a consistent theory in a countable language. If a partial type $\Sigma(x)$ is not isolated in T, then there is a countable model of T which omits $\Sigma(x)$.

Reminder

Recall from *Grondslagen van de Wiskunde*:

Theorem

Suppose T is a consistent theory in a language L and C is a set of constants in L. If for any formula $\psi(x)$ in the language L there is a constant $c \in C$ such that

$$T \models \exists x\, \psi(x) \to \psi(c),$$

then T has a model whose universe consists entirely of interpretations of elements of C.

Proof.

Extend T to a maximally consistent theory and then build a model from the constants in C. □

Omitting types theorem, proof

Omitting types theorem

Let T be a consistent theory in a countable language. If a partial type $\Sigma(x)$ is not isolated in T, then there is a countable model of T which omits $\Sigma(x)$.

Proof.

Let $C = \{c_i \,;\, i \in \mathbb{N}\}$ be a countable collection of fresh constants and L_C be the language L extending with these constants. Let $\{\psi_i(x) : i \in \mathbb{N}\}$ be an enumeration of the formulas with one free variable in the language L_C. We will now inductively create a sequence of sentences $\varphi_0, \varphi_1, \varphi_2, \ldots$. The idea is to apply to previous theorem to $T \cup \{\varphi_0, \varphi_1, \ldots\}$.

If $n = 2i$, we take a fresh constant $c \in C$ (one that does not occur in φ_m with $m < n$) and put

$$\varphi_n = \exists x \psi_i(x) \rightarrow \psi(c).$$

This makes sure we can create a model from the constants in C. □

Omitting types theorem, proof finished

Proof.

If $n = 2i + 1$ we make sure that c_i omits $\Sigma(x)$, as follows. Consider $\delta = \bigwedge_{m<n} \varphi_m$. δ is really of the form $\delta(c_i, \overline{c})$ where $\overline{c}$ is a sequence of constants not containing c_i. Since $\Sigma(x)$ is not isolated, there must be a formula $\sigma(x) \in \Sigma(x)$ such that $T \not\models \exists \overline{y}\, \delta(x, \overline{y}) \to \sigma(x)$; in other words, such that $T \cup \{\exists \overline{y}\, \delta(x, y)\} \cup \{\neg\sigma(x)\}$ is consistent. Put $\varphi_{2n} = \neg\sigma(c_i)$.

The proof is now finished by showing by induction that each $T \cup \{\varphi_0, \ldots, \varphi_n\}$ is consistent and then applying the theorem from *Grondslagen*. □

Exercises

Exercise

Prove the generalised omitting types theorem: Let T be a consistent theory in a countable language and let $\{\Gamma_i : i \in \mathbb{N}\}$ be a sequence of partial n_i-types (for varying n_i). If none of the Γ_i is isolated in T, then there is a countable model which omits all Γ_i.

Exercise

Let T be a complete theory. Show that models of T realise all isolated partial types.

Exercise

Prove that the omitting types theorem is specific to the countable case: give an example of a consistent theory T in an uncountable language and a partial type in T which is not isolated, but which is nevertheless realised in every model of T.

ω-categoricity

Convention

Let us say a theory is *nice* if it

- is complete,
- and formulated in a countable language,
- and has infinite models.

Definition

A theory is *ω-categorical* if all its countably infinite models are isomorphic.

Theorem (Ryll-Nardzewski)

For a nice theory T the following are equivalent:

1. T is ω-categorical;
2. all n-types are isolated;
3. all models of T are ω-saturated;
4. all countable models of T are ω-saturated.

Remark

Note that for any theory T we have:

Proposition

The following are equivalent: (1) all n-types are isolated; (2) every $S_n(T)$ is finite; (3) for every n there are only finite many formulas $\varphi(x_1, \ldots, x_n)$ up to equivalence relative to T.

Proof.

(1) $\Leftrightarrow$ (2) holds because $S_n(T)$ is a compact Hausdorff space.
(2) $\Rightarrow$ (3): If there are only finitely many types, then each of these isolated, so there are formulas $\psi_1(x_1, \ldots, x_n), \ldots, \psi_m(x_1, \ldots, x_n)$ "isolating" all these types with $T \models \bigvee_i \psi_i$. But then every formula $\varphi(x_1, \ldots, x_n)$ is equivalent to the disjunction of the ψ_i of which it is a consequence.
(3) $\Rightarrow$ (2): If every formula $\varphi(x_1, \ldots, x_n)$ is equivalent modulo T to one of $\psi_1(x_1, \ldots, x_n), \ldots, \psi_m(x_1, \ldots, x_n)$, then every n-type is completely determined by saying which ψ_i it does and which it does not contain. □

Ryll-Nardzewski Theorem

Theorem (Ryll-Nardzewski)

For a nice theory T the following are equivalent:

1. T is ω-categorical;
2. all n-types are isolated;
3. all models of T are ω-saturated;
4. all countable models of T are ω-saturated.

Proof.

(1) $\Rightarrow$ (2): If T contains a non-isolated type then there is a model where it is realized and a model where it is not realized (by the Omitting Types Theorem). (2) $\Rightarrow$ (3): If all $n+1$-types are isolated, then every 1-type with n parameters from a model is isolated, hence generated by a single formula. So if such a type is finitely satisfiable in a model, that formula can be satisfied there and then the entire type is realised. (3) $\Rightarrow$ (4) is obvious. (4) $\Rightarrow$ (1): Because elementarily equivalent κ-saturated models of cardinality κ are always isomorphic. □

Existence countable saturated models

Corollary

If A is a model and $a_1, \ldots, a_n$ are elements from A, then $\mathrm{Th}(A)$ is ω-categorical iff $\mathrm{Th}(A, a_1, \ldots, a_n)$ is ω-categorical.

Definition

A theory T is *small* if all $S_n(T)$ are at most countable.

Theorem

A nice theory is small iff it has a countable ω-saturated model.

Proof.

$\Leftarrow$: If T is complete and has a countable ω-saturated model, then every type consistent with T is realized in that model. So there are at most countable many n-types for any n.

$\Rightarrow$ I will do on the next page. □

Proof finished

Theorem

A nice theory is small iff it has a countable ω-saturated model.

Proof.

$\Rightarrow$: We know that a model A can be elementarily embedded in a model B which realizes all types with parameters from A that are finitely satisfied in A. From the proof of that result we see that if A is a countable and there are at most countably many n-types with a finite set of parameters from A, then all of these types can be realized in a *countable* elementary extension B. Building an ω-chain by repeatedly applying this result and then taking the colimit, we see that A can be embedded in a countable ω-saturated elementary extension. So if A is a countable model of T, we obtain the desired result. □

Vaught's Theorem

Theorem (Vaught)

A nice theory cannot have exactly two countable models (up to isomorphism).

Proof.

Let T be a nice theory. Without loss of generality we may assume that T is small (why?) and not ω-categorical. We will now show that T has at least three models.

First of all, there is a countable ω-saturated model A. In addition, there is a non-isolated type p which is omitted in some model B. Of course, it is realized in A by some tuple $\overline{a}$. Since $\mathrm{Th}(A, \overline{a})$ is not ω-categorical (by the corollary from a few slides back), it has a model different from A. Since this model realizes p, it must be different from B as well. □

Exercises

Exercise

Write down a theory with exactly two countable models.

Exercise

Show for every $n > 2$ there is a nice theory having precisely n countable models (up to isomorphism). (Consider $(\mathbb{Q}, P_0, \ldots, P_{n-2}, c_0, c_1, \ldots)$ where the P_i form a partition into dense subsets and the c_i are an increasing sequence of elements of P_0.)

Exercise

Give an example of a complete theory T in an uncountable language which has exactly one countable model but for which not all $S_n(T)$ are finite.

Prime and atomic models

Definition

Let T be a nice theory.

- A model M of T is called *prime* if it can be elementarily embedded into any model of T.
- A model M of T is called *atomic* if it only realises isolated types (or, put differently, omits all non-isolated types) in $S_n(T)$.

Theorem

A model of a nice theory T is prime iff it is countable and atomic.

Proof.

$\Rightarrow$: Because T is nice it has countable models and non-isolated types can be omitted. For $\Leftarrow$ see the next page. $\square$

Proof continued

Theorem

A model of a nice theory T is prime iff it is countable and atomic.

Proof.

$\Leftarrow$: Let A be a countable and atomic model of a nice theory T and M be any other model of T. Let $\{a_1, a_2, \ldots\}$ be an enumeration of A; by induction on n we will construct an increasing sequence of elementary maps $f_n : \{a_1, \ldots, a_n\} \to M$. We start with $f_0 = \emptyset$, which is elementary as A and M are elementarily equivalent. (They are both models of a complete theory T.)

Suppose f_n has been constructed. The type of $a_1, \ldots, a_{n+1}$ in A is isolated, hence generated by a single formula $\varphi(x_1, \ldots, x_{n+1})$. In particular, $A \models \exists x_{n+1}\, \varphi(a_1, \ldots, a_n, x_{n+1})$, and since f_n is elementary, $M \models \exists x_{n+1}\, \varphi(f_n(a_1), \ldots, f_n(a_n), x_{n+1})$. So choose $m \in M$ such that $M \models \varphi(f_n(a_1), \ldots, f_n(a_n), m)$ and put $f(a_{n+1}) = m$. $\square$

Existence prime models

Theorem

All prime models of a nice theory T are isomorphic. In addition, they are strongly ω-homogeneous.

Proof.

By the familiar back-and-forth techniques. (Exercise!) □

Theorem

A nice theory T has a prime model iff the isolated n-types are dense in $S_n(T)$ for all n.

Remark

Let us call a formula $\varphi(\overline{x})$ *complete* in T if it generates an isolated type in $S_n(T)$: that is, it is consistent and for any other formula $\psi(\overline{x})$ we have either $T \models \varphi(\overline{x}) \to \psi(\overline{x})$ or $T \models \varphi(\overline{x}) \to \neg\psi(\overline{x})$. Then n-types are dense iff every consistent formula $\varphi(\overline{x})$ follows from some complete formula.

Existence prime models, proof

Theorem

A nice theory T has a prime model iff the isolated n-types are dense in $S_n(T)$ for all n.

Proof.

$\Rightarrow$: Let A be a prime model of T. Because a consistent formula $\varphi(\bar{x})$ is realised in *all* models of T, it is realized in A as well, by $\bar{a}$ say. Since A is atomic, $\varphi(\bar{x})$ belongs to the isolated type $\mathrm{tp}_A(\bar{a})$.
$\Leftarrow$: Note that a structure A is atomic iff the sets

$$\Sigma_n(x_1, \ldots, x_n) = \{ \neg\varphi(x_1, \ldots, x_n) : \varphi \text{ is complete} \}$$

are omitted in A. So it suffices to show that the Σ_n are not isolated (by the generalised omitting types theorem). But that holds iff for any consistent $\psi(\bar{x})$ there is a complete formula $\varphi(\bar{x})$ such that $T \not\models \psi(\bar{x}) \to \neg\varphi(\bar{x})$. As $\varphi(\bar{x})$ is complete, this is equivalent to $T \models \varphi(\bar{x}) \to \psi(x)$. So the Σ_n are not isolated iff isolated types are dense. □

Binary trees of formulas

Definition

Let $\{0,1\}^*$ be the set of finite sequences consisting of zeros and ones. A *binary tree* of formulas in variables $\overline{x} = x_1, \ldots, x_n$ (in T) is a collection $\{\varphi_s(\overline{x}) : s \in \{0,1\}^*\}$ such that

- $T \models \big(\varphi_{s0}(\overline{x}) \vee \varphi_{s1}(\overline{x})) \to \varphi_s(\overline{x})\big)$.
- $T \models \neg\big(\varphi_{s0}(\overline{x}) \wedge \varphi_{s1}(\overline{x})\big)$.

Theorem

The following are equivalent for a nice theory T:

(1) $|S_n(T)| < 2^\omega$.

(2) There is no binary tree of consistent formulas in $x_1, \ldots, x_n$.

(3) $|S_n(T)| \leq \omega$.

Clearly, if $\{\varphi_s(\overline{x}) : s \in \{0,1\}^*\}$ is a binary tree of consistent formulas, $\{\varphi_s : s \subseteq \alpha\}$ is consistent for every $\alpha : \mathbb{N} \to \{0,1\}$. This shows (1) $\Rightarrow$ (2). As (3) $\Rightarrow$ (1) is obvious, it remains to show (2) $\Rightarrow$ (3).

A lemma

Lemma

Let T be a nice theory. If $|S_n(T)| > \omega$, then there is a binary tree of consistent formulas in $x_1, \ldots, x_n$.

Proof.

Suppose $|S_n(T)| > \omega$. This implies, since the language of T is countable, that there is a formula $\varphi(\overline{x})$ such that $|[\varphi]| > \omega$. The lemma will now follow from the following *claim*: If $|[\varphi]| > \omega$, then there is a formula $\psi(\overline{x})$ such that $|[\varphi \wedge \psi]| > \omega$ and $|[\varphi \wedge \neg\psi]| > \omega$. Suppose not.
Then $p(\overline{x}) = \{\psi(\overline{x}) \ : \ |[\varphi \wedge \psi]| > \omega\}$ contains a formula $\psi(\overline{x})$ or its negation, but not both, and is closed under logical consequence: so it is a complete type. If $\psi \notin p$, then $|[\varphi \wedge \psi]| \leq \omega$. In addition, the language is countable, so

$$[\varphi] = \bigcup_{\psi \notin p} [\varphi \wedge \psi] \cup \{p\}$$

is a countable union of countable sets and hence countable, contradicting our choice of φ. $\square$

Small theories have prime models

Corollary

If T is nice and $|S_n(T)| < 2^\omega$ for all n, then T is small.

Corollary

If T is nice and small, then isolated types are dense. So T has a prime model.

Proof.

If isolated types are not dense, then there is a consistent $\varphi(\bar{x})$ which is not a consequence of a complete formula. Call such a formula *perfect*. Since perfect formulas are not complete, they can be "decomposed" into two consistent formulas which are jointly inconsistent. These have to be perfect as well, leading to a binary tree of consistent formulas. □

Stability

Let κ be an infinite cardinal.

Definition

A theory T is *κ-stable* if in each model of T, over set of parameters of size at most κ, and for each n, there are at most κ many n-types. That is:

$$|A| \leq \kappa \Rightarrow |S_n(A)| \leq \kappa.$$

An easy induction argument shows that it suffices to require that $|A| \leq \kappa \Rightarrow |S_1(A)| \leq \kappa$.

The theory ACF_0 is ω-stable, but DLO and $RCOF$ are not!

Goal of the day

Theorem

A countable theory T which is categorical in an uncountable cardinal is ω-stable.

By the way, by a countable theory I mean a theory in a countable language. For the proof I need two ingredients:

1. Ramsey's Theorem: a result from combinatorics.
2. The notion of (order) indiscernible.

Ramsey's Theorem

Ramsey's Theorem

Let A be infinite and $n \in \mathbb{N}$. Partition $[A]^n$, the set of n-element subsets of A, into subsets $C_1, \ldots, C_k$ (their *colours*). Then there is an infinite subset of A all whose n-element subsets belong to the same subset C_i.

Proof.

By induction on n. $n = 1$ is the pigeon hole principle. So we assume the statement is true for n and prove it for $n+1$. Let $a_0 \in A$: then any colouring of $[A]^{n+1}$ induces a colouring of $[A \setminus \{a_0\}]^n$: just colour $\alpha \in [A \setminus \{a_0\}]$ by the colour of $\{a_0\} \cup \alpha$. We obtain a infinite monochromatic subset $B_1 \subseteq A \setminus \{a_0\}$. Picking an element $a_1 \in B_1$ and continuing in this fashion we obtain an infinitely descending sequence $A = B_0 \supseteq B_1 \supseteq \ldots$ and elements $a_i \in B_i - B_{i+1}$ such that the colour of any $(n+1)$-element subset $\{a_{i(0)}, \ldots, a_{i(n)}\}$ $(i(0) < \ldots < i(n))$ depends only on the value of $i(0)$. By the pigeon hole principle there are infinitely many $i(0)$ for which this colour will be the same. These $a_{i(0)}$ then yield the desired monochromatic set. □

Indiscernibles

Definition

Let I be a linear order and A be an L-structure. A family of elements $(a_i)_{i \in I}$ (or tuples of elements, all of the same length) is called a *sequence of indiscernibles* if for all formulas $\varphi(x_1, \ldots, x_n)$ and all $i_1 < \ldots < i_n$ and $j_1 < \ldots, < j_n$ from I we have

$$A \models \varphi(i_1, \ldots, i_n) \leftrightarrow \varphi(j_1, \ldots, j_n).$$

Definition

Let I be an infinite linear order and $\mathcal{I} = (a_i)_{i \in I}$ be a sequence of elements in M, $A \subseteq M$. The *Ehrenfeucht-Mostowski type* $\mathrm{EM}(\mathcal{I}/A)$ of $\mathcal{I}$ over A is the set of $L(A)$-formulas $\varphi(x_1, \ldots, x_n)$ with $M \models \varphi(a_{i_1}, \ldots, a_{i_n})$ for all $i_1 < \ldots < i_n$.

Note that if $(a_i)_{i \in I}$ is a sequence of indiscernibles, then the Ehrenfeucht-Mostowski type $\mathrm{EM}(\mathcal{I}/A)$ is complete (contains either a formula or its negation).

The Standard Lemma

The Standard Lemma

Let I and J be two infinite linear orders and $\mathcal{I} = (a_i)_{i \in I}$ be a sequence of distinct elements of a structure M. Then there is a structure $N \equiv M$ with an indiscernible sequence $(b_j)_{j \in J}$ realizing the Ehrenfeucht-Mostowski type $\mathrm{EM}(\mathcal{I}/A)$.

Proof.

Choose a set C of new constants with an ordering isomorphic to J. We need to show that

$$\mathrm{Th}(M) \cup \{\varphi(\overline{c}) \,:\, \varphi(\overline{x}) \in \mathrm{EM}(\mathcal{I}/A)\} \cup \{\varphi(\overline{c}) \leftrightarrow \varphi(\overline{d}) \,:\, \overline{c}, \overline{d} \in C\}$$

is consistent. (Here the $\varphi(\overline{x})$ are L-formulas and $\overline{c}, \overline{d}$ tuples in increasing order.) □

Proof of The Standard Lemma, finished

Proof.

By compactness it is sufficient to show that

$$\mathrm{Th}(M) \cup \{\varphi(\overline{c}) : \varphi(\overline{x}) \in \mathrm{EM}(\mathcal{I}/A), \overline{c} \in C_0\} \cup \\ \{\varphi(\overline{c}) \leftrightarrow \varphi(\overline{d}) : \varphi(\overline{x}) \in \Delta, \overline{c}, \overline{d} \in C_0\}$$

has a model, where C_0 and Δ are finite. In addition, we may assume that all tuples $\overline{c}$ have the same length n.

In that case we may define an equivalence relation $\sim$ on $[A]^n$ by

$$\overline{a} \sim \overline{b} \Leftrightarrow M \models \varphi(\overline{a}) \leftrightarrow \varphi(\overline{b}) \text{ for all } \varphi(x_1, \ldots, x_n) \in \Delta$$

where $\overline{a}, \overline{b}$ are tuples in increasing order. Since this equivalence relation has at most $2^{|\Delta|}$ equivalence classes, there is an infinite subset B of A with all n-elements subsets in the same equivalence class. Interpret $c \in C_0$ by elements b_c in B ordered in the same way as the c. Then $(M, b_c)_{c \in C_0}$ is a model. □

Another lemma

Corollary

Assume T has an infinite model. Then, for any linear order I, the theory T has a model with a sequence $(a_i)_{i \in I}$ of distinct indiscernibles.

Lemma

Assume L is countable. If the L-structure M is generated by a well-ordered sequence $(a_i)_{i \in I}$ of indiscernibles, then M realises only countably many types over every countable subset of M.

Proof.

See handout. □

Another corollary

Corollary

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

Proof.

Let T' be the skolemisation of T in richer language $L' \supseteq L$, and let I be a well-ordering of cardinality κ and N' be a model of T' with indiscernibles $(a_i)_{i \in I}$. Then the Skolem hull M' generated by $(a_i)_{i \in I}$ has cardinality κ and is an elementary substructure of N'. In addition, it realises only countably many types over every countable subset by the previous lemma. But then the same is certainly also true for the reduct $M = M' \restriction L$. □

Goal of the day achieved

Theorem

A countable theory T which is categorical in an uncountable cardinal is ω-stable.

Proof.

Let N be a model and $A \subseteq N$ countable with $S(A)$ uncountable. Let $(b_i)_{i \in I}$ be a sequence of ω_1-many elements realizing different types over A. First choose an elementary substructure M_0 of N of cardinality ω_1 which contains both A and the b_i, and then choose an elementary extension M of M_0 of cardinality κ. The model M is of cardinality κ and realises uncountably many types over the countable set A. But by the previous corollary T also has a model of cardinality κ in which this is not the case. So T is not κ-categorical. □

Next goals

The next step in the proof of Morley's Theorem is an analysis of nice ω-stable theories. In particular, we need to establish the following three results for such theories T:

Theorem

T is κ-stable for all $\kappa \geq \omega$.

Theorem

Suppose $A \models T$ and $C \subseteq A$, where A is uncountable and $|C| < |A|$. Then there exists a sequence of distinct indiscernibles in $(A, a)_{a \in C}$.

Theorem

Suppose $A \models T$ and $C \subseteq A$. There exists $B \preceq A$ such that $C \subseteq B$ and B is atomic over C.

To prove these results we need the notions of *Morley rank* and *Morley degree*.

Definition of $\mathrm{RM} \geq \alpha$

Today we will fix a complete theory T.

Definition

Suppose $A \models T$, $\varphi(x)$ is an L_A-formula, and α is an ordinal. We define $\mathrm{RM}_x(A, \varphi(x)) \geq \alpha$ by induction on α:

1. $\mathrm{RM}_x(A, \varphi(x)) \geq 0$ if $A \models \exists x\, \varphi(x)$;
2. $\mathrm{RM}_x(A, \varphi(x)) \geq \alpha + 1$ if there is an elementary extension B of A and a sequence $(\varphi_k(x) : k \in \mathbb{N})$ of L_B-formulas such that
   1. $B \models \forall x\, (\, \varphi_k(x) \to \varphi(x)\,)$ for all $k \in \mathbb{N}$;
   2. $B \models \forall x\, \neg(\, \varphi_k(x) \wedge \varphi_l(x)\,)$ for all distinct $k, l \in \mathbb{N}$;
   3. $\mathrm{RM}_x(B, \varphi_k(x)) \geq \alpha$ for all $k \in \mathbb{N}$;
3. for λ a limit ordinal, $\mathrm{RM}_x(A, \varphi(x)) \geq \lambda$ if $\mathrm{RM}_x(A, \varphi(x)) \geq \alpha$ for all $\alpha < \lambda$.

Main property of $\mathrm{RM} \geq \alpha$

Lemma

Suppose $A \models T$ and $\varphi(x)$ is an L_A-formula. Let S be the set of ordinals α such that $\mathrm{RM}_x(A, \varphi(x)) \geq \alpha$ holds. Then exactly one of the following alternatives holds:

1. S is empty;
2. S is the class of all ordinals;
3. $S = \{\alpha : \alpha \leq \gamma\}$ for some ordinal γ.

Proof.

This really amounts to showing that $\mathrm{RM}_x(A, \varphi(x)) \geq \alpha$ and $\alpha > \beta \geq 0$ imply $\mathrm{RM}_x(A, \varphi(x)) \geq \beta$. We prove this by induction on α and β. The cases where α or β is a limit ordinal are easy, so assume $\mathrm{RM}_x(A, \varphi(x)) \geq \alpha + 1$ and $\alpha + 1 > \beta + 1$ (so $\alpha > \beta$). The first assumption implies that there is an elementary extension B of A and a sequence $(\varphi_k(x) : k \in \mathbb{N})$ with $\mathrm{RM}_x(B, \varphi_k(x)) \geq \alpha$. But then $\mathrm{RM}_x(B, \varphi_k(x)) \geq \beta$ and hence $\mathrm{RM}_x(A, \varphi(x)) \geq \beta + 1$, as desired. □

Morley rank

Definition

Let A be a model of T and let $\varphi(x)$ be an L_A-formula. $\mathrm{RM}_x(A, \varphi(x)) \geq \alpha$ is false for all ordinals α, then we write $\mathrm{RM}_x(A, \varphi(x)) = -\infty$. If $\mathrm{RM}_x(A, \varphi(x)) \geq \alpha$ holds for all ordinals α, then we write $\mathrm{RM}_x(A, \varphi(x)) = +\infty$. Otherwise we define $\mathrm{RM}_x(A, \varphi(x))$ to be the greatest ordinal α for which $\mathrm{RM}_x(A, \varphi(x)) \geq \alpha$ holds, and we say that $\varphi(x)$ is *ranked*.

Morley rank depends on the type only

Lemma

Let A be a model and $\varphi(x, y)$ be an L-formula. If a is a finite tuple of elements of A, then the value of $\mathrm{RM}_x(A, \varphi(x, a))$ depends only on $\mathrm{tp}_A(a)$.

Proof.

It suffices to prove that the truth value of $\mathrm{RM}_x(A, \varphi(x, a)) \geq \alpha$ only depends on the type of a. We prove this by induction on α; the case that $\alpha = 0$ or a limit ordinal is trivial. So assume the statement holds for all $\alpha < \beta + 1$.

For $j = 1, 2$, let A_j be a model of T and a_j be a finite tuples from A_j with $\mathrm{tp}_{A_1}(a_1) = \mathrm{tp}_{A_2}(a_2)$. We assume $\mathrm{RM}_x(A_1, \varphi(x, a_1)) \geq \beta + 1$ and need to prove $\mathrm{RM}_x(A_2, \varphi(x, a_2)) \geq \beta + 1$.

The assumption yields an elementary extension B_1 of A_1 and a sequence of formulas $(\varphi_k(x, b_k) : k \in \mathbb{N})$ to witness that $\mathrm{RM}_x(A_1, \varphi(x, a_1)) \geq \beta + 1$, that is, ...

Morley rank depends on the type only, continued

Proof.

1. $B_1 \models \forall x\,(\,\varphi_k(x, b_k) \to \varphi(x, a_1)\,)$ for all $k \in \mathbb{N}$;
2. $B_1 \models \forall x\,\neg(\,\varphi_k(x, b_k) \wedge \varphi_l(x, b_l)\,)$ for all distinct $k, l \in \mathbb{N}$;
3. $\mathrm{RM}_x(B_1, \varphi_k(x, b_k)) \geq \beta$ for all $k \in \mathbb{N}$.

Now let B_2 be *any* ω-saturated elementary extension of A_2. We know that $\mathrm{tp}_{B_1}(a_1) = \mathrm{tp}_{B_2}(a_2)$. Since B_2 is ω-saturated, we may construct inductively a sequence $(c_k\,:\,k \in \mathbb{N})$ of finite tuples from B_2 such that for all $k \in \mathbb{N}$

$$\mathrm{tp}_{B_2}(a_2 c_0 \ldots c_k) = \mathrm{tp}_{B_1}(a_1 b_0 \ldots b_k).$$

It follows that

1. $B_2 \models \forall x\,(\,\varphi_k(x, c_k) \to \varphi(x, a_2)\,)$ for all $k \in \mathbb{N}$;
2. $B_1 \models \forall x\,\neg(\,\varphi_k(x, c_k) \wedge \varphi_l(x, c_l)\,)$ for all distinct $k, l \in \mathbb{N}$;
3. $\mathrm{RM}_x(B_2, \varphi_k(x, c_k)) \geq \beta$ for all $k \in \mathbb{N}$.

(Statements (1) and (2) are immediate; for (3) use the induction hypothesis.) So $\mathrm{RM}_x(B_2, \varphi_k(x, a_2)) \geq \beta + 1$. □

Exercises

Exercise

Let A be an ω-saturated model of T and let $\varphi(x)$ be an L_A-formula. In applying the definition of $\mathrm{RM}_x(A, \varphi(x)) \geq \alpha$ one may take the elementary extension B to be A itself.

Exercise (Properties of Morley rank)

Let A be a model of T and let $\varphi(x), \psi(x)$ be L_A-formulas.

1. $\mathrm{RM}_x(A, \varphi(x)) = 0$ iff the number of tuples $u \in A$ for which $A \models \varphi(u)$ is finite and > 0.
2. if $A \models \varphi(x) \to \psi(x)$, then $\mathrm{RM}_x(A, \varphi(x)) \leq \mathrm{RM}_x(A, \psi(x))$.
3. $\mathrm{RM}_x(A, \varphi(x) \vee \psi(x)) = \max(\mathrm{RM}_x(A, \varphi(x)), \mathrm{RM}_x(A, \psi(x)))$.
4. if $\varphi(x)$ is ranked and $\mathrm{RM}_x(A, \varphi(x)) > \beta$, then there exists an elementary extension B of A and an L_B-formula $\chi(x)$ such that $B \models \chi(x) \to \varphi(x)$ and $\mathrm{RM}_x(B, \chi(x)) = \beta$.

Towards Morley degree

Lemma

Let A be a model of T and $\varphi(x)$ be a ranked L_A-formula. There exists a finite bound on the integers k such that there exists an elementary extension B of A and L_B-formulas $(\varphi_j(x) : 0 \leq j < k)$ such that

1. $\mathrm{RM}_x(B, \varphi_j(x)) = \mathrm{RM}_x(A, \varphi(x))$ for all $j < k$;
2. $B \models (\varphi_j(x) \rightarrow \varphi(x))$ for all $j < k$;
3. $B \models \neg(\varphi_i(x) \wedge \varphi_j(x))$ for distinct $i, j < k$.

Moreover, the maximum value of k depends only on $\mathrm{tp}_A(a)$. And if A is ω-saturated, a maximal sequence can be found for B equal to A itself.

Proof. Write $\varphi(x) = \varphi(x, a)$ where $\varphi(x, y)$ is an L-formula. The existence of an elementary extension B and L_B-formulas $\varphi_j(x)$ having properties (1)-(3) amounts to the consistency of a certain set of sentences involving a and the parameters from B occurring in the $\varphi_j(x)$. So consistency depends solely on the type of a; and these sentences will be realized in any ω-saturated extension of A, if consistent.

Towards Morley degree, continued

Proof.

So we may assume that A is ω-saturated and restrict ourselves to considering sequences of L_A-formulas $(\varphi_j(x) : 0 \leq j < k)$.

We will create a binary tree of L_A-formulas, each having Morley rank α. We put $\varphi_{<>} = \varphi(x)$. If φ_σ has been constructed, we check whether there is a formula ψ such that both $\varphi \wedge \psi$ and $\varphi \wedge \neg\psi$ have Morley rank α. If so, we put $\varphi_{\sigma 0} = \varphi \wedge \psi$ and $\varphi_{\sigma 1} = \varphi \wedge \neg\psi$ for some such ψ. Otherwise we stop.

The resulting tree has to be finite: for otherwise it would have (by König's Lemma) an infinite branch α. But then $\varphi_{\overline{\alpha}(n)} \wedge \neg\varphi_{\overline{\alpha}(n+1)}$ would be an infinite sequence witnessing that the Morley rank of φ is $\geq \alpha + 1$.

Let L be the collection of leaves of the tree. Then $(\varphi_s : s \in L)$ is a sequence satisfying (1)-(3): in fact, $\varphi \leftrightarrow \bigvee_{s \in L} \varphi_s$. We claim it is maximal. □

Towards Morley degree, finished

Proof.

For suppose $(\psi_j(x) : 0 \leq j < k)$ is another such sequence satisfying (1)-(3) and $k > |S_0|$. Since $\psi_i(x)$ and $\psi_j(x)$ are contradictory whenever i and j are distinct, at most one of $\varphi_s \wedge \psi_i$ and $\varphi_s \wedge \psi_j$ can have Morley rank α. Since $k > |S_0|$, it follows from the pigeonhole principle that there is a $j < k$ such that $\psi_j \wedge \varphi_s$ has rank $< \alpha$ for all $s \in S_0$. But as ψ_j is equivalent to the disjunction of all formulas $\psi_j \wedge \varphi_s$, it follows that ψ_j must itself have Morley rank $< \alpha$. Contradiction! □

Definition

Given a ranked L_A-formula $\varphi(x)$, the greatest integer whose existence we just proved is called the *Morley degree* of $\varphi(x)$ and it is denoted by $dM(\varphi(x))$.

Properties of Morley degree

Lemma

Let A be an ω-saturated model of T and let $\varphi(x)$ and $\psi(x)$ be ranked L_A-formulas.

1. If $dM(\varphi(x)) = d$ and this is witnessed by the sequence $(\varphi_j(x) : 0 \leq j < d)$, then each $\varphi_j(x)$ has Morley degree 1.
2. If $\mathrm{RM}_x(A, \varphi(x)) = \mathrm{RM}_x(A, \psi(x))$ and $A \models \varphi(x) \to \psi(x)$, then $dM(\varphi(x)) \leq dM(\psi(x))$.
3. If $\mathrm{RM}_x(A, \varphi(x)) = \mathrm{RM}_x(A, \psi(x))$, then $dM(\varphi(x) \vee \psi(x)) \leq dM(\varphi(x)) + dM(\psi(x))$, with equality if $A \models \neg(\,\varphi(x) \wedge \psi(x)\,)$.
4. If $\mathrm{RM}_x(A, \varphi(x)) < \mathrm{RM}_x(A, \psi(x))$, then $dM(\varphi(x) \vee \psi(x)) = dM(\varphi(x))$.

Proof.

Exercise! □

Types and Morley rank

Lemma

Let $A \models T$ and $C \subseteq A$. Let $p(x)$ be a type in L_C that is consistent with $\mathrm{Th}((A, a)_{a \in C})$. Assume that some formula in $p(x)$ is ranked. Then there exists a formula $\varphi_p(x)$ in $p(x)$ that determines $p(x)$ in the following sense:

> *$p(x)$ consists exactly of the L_C-formulas $\psi(x)$ such that $\mathrm{RM}(\psi(x) \wedge \varphi_p(x)) = \mathrm{RM}(\varphi_p(x))$ and $dM(\psi(x) \wedge \varphi_p(x)) = dM(\varphi_p(x))$.*

Indeed, such a formula can be obtained by taking $\varphi_p(x)$ to be a formula $\varphi(x)$ in $p(x)$ with least possible Morley rank and Morley degree, in lexicographic order.

Proof.

Choose $\varphi_p(x)$ as in the last sentence of the lemma. Then, if $\psi(x)$ is any formula in $p(x)$, also $\psi(x) \wedge \varphi_p(x) \in p(x)$ and hence $\mathrm{RM}(\psi(x) \wedge \varphi_p(x)) \geq \mathrm{RM}(\varphi_p(x))$ by choice of $\varphi_p(x)$. Hence $\mathrm{RM}(\psi(x) \wedge \varphi_p(x)) = \mathrm{RM}(\varphi_p(x))$. Similarly for Morley degree. □

Types and Morley rank, continued

Proof.

Conversely, suppose $\psi(x)$ is any L_C-formula with $\mathrm{RM}(\psi(x) \wedge \varphi_p(x)) = \mathrm{RM}(\varphi_p(x))$ and $dM(\psi(x) \wedge \varphi_p(x)) = dM(\varphi_p(x))$. By way of contradiction, if $\psi(x) \notin p(x)$, then $\neg\psi(x) \in p(x)$. But then $\mathrm{RM}(\neg\psi(x) \wedge \varphi_p(x)) = \mathrm{RM}(\varphi_p(x))$, in which case we have $dM(\varphi_p(x)) \geq dM(\psi(x) \wedge \varphi_p(x)) + dM(\neg\psi(x) \wedge \varphi_p(x)) > dM(\psi(x) \wedge \varphi_p(x))$, which is a contradiction. □

Definition

Let $p(x)$ be a type as in the statement of the lemma. Then we define $\mathrm{RM}(p(x))$ to be the least Morley rank of a formula in $p(x)$. If some formula in $p(x)$ is ranked, we define $dM(p(x))$ to be the least Morley degree of a formula $\varphi(x)$ in $p(x)$ that satisfies $\mathrm{RM}(\varphi(x)) = \mathrm{RM}(p(x))$.

Totally transcendental theories

Definition

A theory T is *totally transcendental* if it has no model M with a binary tree of consistent $L(M)$-formulas.

Theorem

Let L be countable. Then the following conditions are equivalent:

1. T is ω-stable;
2. T is totally transcendental;
3. if $A \models T$ and $\varphi(x)$ is an L_A-formula which is realized in A, then $\varphi(x)$ is ranked;
4. T is λ-stable for all $\lambda \geq \omega$.

Proof.

(1) $\Rightarrow$ (2): In a binary tree of consistent $L(M)$-formulas only countably many parameters from M occur; but its existence implies that there are at least 2^ω different types over this countable set. □

Proof continued

Proof.

(2) $\Rightarrow$ (3): Let M be an ω-saturated model of T and let $\varphi(x)$ be a formula of Morley rank $+\infty$. Since the formulas from L_M form a set, there is an ordinal α such that any formula $\psi(x)$ whose Morley rank is $\geq \alpha$ has Morley rank is $+\infty$. So because $\mathrm{RM}(\varphi(x)) \geq \alpha + 1$, there must be contradictory formulas $\psi_1(x)$ and $\psi_2(x)$ with $\mathrm{RM}(\psi_i(x)) \geq \alpha$ and $M \models \psi_i(x) \rightarrow \varphi(x)$. So $\varphi(x) \wedge \psi_1(x)$ and $\varphi(x) \wedge \psi_2(x)$ both have Morley rank $+\infty$. Continuing in this way we create a binary tree of consistent formulas in M.

(3) $\Rightarrow$ (4): Let $A \models T$ and $C \subseteq A$ with $|C| \leq \lambda$. Then every type $p(x)$ is uniquely determined by an L_C-formula $\varphi_p(x)$. Since there are at most λ many L_C-formulas (L is countable!), there are at most λ many types.

(4) $\Rightarrow$ (1) is obvious. □

Second theorem

Today all theories are assumed to be *nice*.

Notation

Let A be an L-structure. If b is a tuple in A and B is any subset of A, we will write $\mathrm{tp}_A(b/B)$ for the type in L_B realized by b.

Theorem

Assume T is an ω-stable theory, and suppose $A \models T$ and $C \subseteq A$. If A is uncountable and $|C| < |A|$, then there is a nonconstant sequence of indiscernibles in $(A, a)_{a \in C}$.

Proof.

We may assume C is infinite. Write $\lambda = |C|$. The formula $x = x$ is satisfied by $> \lambda$ many elements, so choose an L_A-formula $\varphi(x)$ that is satisfied by $> \lambda$ many elements and has minimum possible Morley rank and degree; say these are (α, d). Note that $\alpha > 0$ since $\varphi(x)$ is satisfied by infinitely many elements. By adding finitely many elements to C we may assume that $\varphi(x)$ is an L_C-formula. □

Second theorem, proof continued

Proof.

We will construct a sequence $(a_k : k \in \mathbb{N})$ of elements of A that satisfy $\varphi(x)$ and such that Morley rank and degree of $\mathrm{tp}_A(a_k/C \cup \{a_0, \ldots, a_{k-1}\})$ is exactly (α, d).

First we claim that there is an a_0 with this property. For if no such element would exist, we would have that Morley rank and degree of $\mathrm{tp}_A(a/C)$ is $< (\alpha, d)$ for all $a \in A$ satisfying $\varphi(x)$. So each $a \in A$ which satisfies $\varphi(x)$ also satisfies an L_C-formula $\psi_a(x)$ with Morley degree and rank $< (\alpha, d)$. But since there are at most λ many L_C-formulas and more than λ many a satisfying $\varphi(x)$, there must be a formula with Morley rank and degree $< (\alpha, d)$ satisfied by $> \lambda$ many a. Contradiction! The construction of a_k given $a_0, \ldots, a_{k-1}$ is similar. So the result follows from the following technical lemma. □

Technical lemma

Lemma

Assume T is ω-stable and suppose $A \models T$ and $C \subseteq A$. Let $\varphi(x)$ be a ranked L_C-formula, and set $(\alpha, d) = (\mathrm{RM}(\varphi(x)), dM(\varphi(x)))$. Suppose $(a_k : k \in \mathbb{N})$ is a sequence of tuples and write $p_k(x) = \mathrm{tp}_A(a_k / C \cup \{a_0, \ldots, a_{k-1}\})$. If $A \models \varphi(a_k)$ and $(\mathrm{RM}(p_k(x)), dM(p_k(x))) = (\alpha, d)$, then $(a_k : k \in \mathbb{N})$ is an indiscernible sequence in $(A, a)_{a \in C}$.

Proof.

Exercise! Hint: Prove by induction on n that whenever $i_0 < \ldots < i_n$, then $\mathrm{tp}(a_{i_0}, \ldots, a_{i_n}/C) = \mathrm{tp}(a_0, \ldots, a_n/C)$ and use the lemma on types and Morley rank and degree. □

Third goal

Recall that the third goal was:

Theorem

Assume T is ω-stable. If $A \models T$ and $C \subseteq A$, then there exists $B \preceq A$ such that $C \subseteq B$ and B is atomic over C.

We do this in two steps: first we show that we can find such a B where B is *constructible* over C; and then we show that constructible extensions have to be atomic.

Definition

Let A be an L-structure and $C \subseteq A$. We say that A *is constructible over* C if there is an ordinal γ and an enumeration $A = (a_\alpha : \alpha < \gamma)$ such that each a_α is atomic over $C \cup A_\alpha$, where $A_\alpha = \{a_\mu : \mu < \alpha\}$.

Existence constructible extensions

Theorem

Assume T is ω-stable. If $A \models T$ and $C \subseteq A$, then there exists $B \preceq A$ such that $C \subseteq B$ and B is constructible over C.

Proof.

T is totally transcendental, so if B is a subset of a model A of T, then $\mathrm{Th}(A_B)$ has no binary tree of consistent formulas. So isolated types in $\mathrm{Th}(A_B)$ are dense.

Now use Zorn's Lemma to find a maximal construction $(a_\alpha)_{a<\lambda}$ which cannot be prolonged by an element $a_\lambda \in M$. Clearly C is contained in A_λ. We show that A_λ is the universe of an elementary substructure by using the Tarski-Vaught Test. So assume $\varphi(x)$ is an L_{A_λ}-formula and $A \models \exists x\, \varphi(x)$. Since isolated types over A_λ are dense, there is an isolated $p(x) \in S(A_\lambda)$ with $\varphi(x) \in p(x)$. Let b be a realisation of $p(x)$ in A. If $b \notin A_\lambda$, then we could prolong our construction by $a_\lambda = b$; thus $b \in A_\lambda$ and $\varphi(x)$ is realised in A_λ. □

Useful lemma

Lemma

Let a and b be two finite tuples of elements of a structure M. Then $\mathrm{tp}(ab)$ is atomic if and only if $\mathrm{tp}(a/b)$ and $\mathrm{tp}(b)$ are atomic.

Proof.

First assume that $\varphi(x, y)$ isolates $\mathrm{tp}(a, b)$. Then $\varphi(x, b)$ isolates $\mathrm{tp}(a/b)$ and we claim $\exists x\, \varphi(x, y)$ isolates $p(y) = \mathrm{tp}(b)$: we have $\exists x\, \varphi(x, y) \in p(y)$ and if $\sigma(y) \in p(y)$, then $M \models \forall x, y\, (\, \varphi(x, y) \to \sigma(y)\,)$ and hence $M \models \forall y\, (\, \exists x\, \varphi(x, y) \to \sigma(y)\,)$.

Conversely, suppose $\rho(x, b)$ isolates $\mathrm{tp}(a/b)$ and $\sigma(y)$ isolates $p(y) = \mathrm{tp}(b)$. Then $\rho(x, y) \wedge \sigma(y)$ isolates $\mathrm{tp}(a, b)$. For if $\varphi(x, y) \in \mathrm{tp}(a, b)$, then $\varphi(x, b)$ belongs to $\mathrm{tp}(a/b)$ and $M \models \forall x\, (\, \rho(x, b) \to \varphi(x, b)\,)$. Hence $\forall x\, (\, \rho(x, y) \to \varphi(x, y)) \in p(y)$ and so it follows that $M \models \forall y\, (\, \sigma(y) \to \forall x\, (\, \rho(x, y) \to \varphi(x, y)\,)$. Thus $M \models \forall x, y\, (\, \rho(x, y) \wedge \sigma(y) \to \varphi(x, y)\,)$. $\square$

Constructible extensions are atomic

Lemma

Constructible extensions are atomic.

Proof.

Let M_0 be a constructible extension of A and let $\overline{a}$ be a tuple from M_0. We have to show that $\overline{a}$ is atomic over A. We can clearly assume that the elements of $\overline{a}$ are pairwise distinct and do not belong to A. We can permute the elements of $\overline{a}$ so that

$$\overline{a} = a_\alpha \overline{b}$$

for some tuple $\overline{b} \in A_\alpha$. Let $\varphi(x, \overline{c})$ be an $L(A_\alpha)$-formula which is complete over A_α and satisfied by a_α. The a_α is also atomic over $A \cup \{\overline{b}\overline{c}\}$. Using induction, we know that $\overline{b}\overline{c}$ is atomic over A. So by the previous lemma $a_\alpha \overline{b}\overline{c}$ and $\overline{a} = a_\alpha \overline{b}$ are atomic over A. □

κ-categoricity and saturation

Theorem

A theory T is κ-categorical if and only if all models of cardinality κ are κ-saturated.

For the proof we need a lemma:

Lemma

If T is κ-stable, then for all regular $\lambda \leq \kappa$ there is a model of cardinality κ which is λ-saturated.

Proof.

We constuct a sequence $(M_\alpha : \alpha \in \lambda)$ of models of T of cardinality κ: we start with any model M_0 of cardinality κ of T; at limit stages we take the colimit and at successor stages we take a model $M_{\alpha+1}$ which realises all types in $S(M_\alpha)$. This we can do with a model of cardinality κ since $|S(M_\alpha)| \leq \kappa$. The colimit of the entire chain will be λ-saturated. □

κ-categoricity and saturation: proof

Theorem

A theory T is κ-categorical if and only if all models of cardinality κ are κ-saturated.

Proof.

Note that we already proved this result for $\kappa = \omega$ and that we also know that any two κ-saturated models of cardinality κ are isomorphic. So we only need to show that if T is κ-categorical for some uncountable cardinal κ, then all models of cardinality κ are κ-saturated.

But then T is ω-stable, hence totally transcendental, hence κ-stable. So by the lemma the unique model of T of cardinality κ is μ^+-saturated for all $\mu < \kappa$. So this model is κ-saturated. □

A theorem implying Morley's theorem

So Morley's Theorem will follow from:

Theorem

Suppose T is ω-stable and assume κ is an uncountable cardinal and that every model of T of cardinality κ is κ-saturated. Then every uncountable model of T is saturated.

Proof.

Suppose T is ω-stable and T has a model of cardinality λ that is not λ-saturated. (Goal is to construct a model of cardinality κ that is not κ-saturated.) So there is a subset C of A of cardinality $< \lambda$ and a type $p(x)$ over C such that $p(x)$ is consistent with $\mathrm{Th}((A, a)_{a \in C})$ but not realized in $(A, a)_{a \in C}$. We know that there is a nonconstant sequence $(a_k : k \in \mathbb{N})$ of indiscernibles in $(A, a)_{a \in C}$ (second goal). Write $I = \{a_k : k \in \mathbb{N}\}$ and note that (*): *for each $L(C \cup I)$-formula $\varphi(x)$ that is satisfiable in $(A, a)_{a \in C \cup I}$ there exists $\psi(x) \in p(x)$ such that $\varphi(x) \wedge \neg\psi(x)$ is satisfiable in $(A, a)_{a \in C \cup I}$.* (For otherwise $p(x)$ would be realized in $(A, a)_{a \in C}$.) □

A theorem implying Morley's theorem, proof continued

Proof.

We have (*): *for each $L(C \cup I)$-formula $\varphi(x)$ that is satisfiable in $(A, a)_{a \in C \cup I}$ there exists $\psi(x) \in p(x)$ such that $\varphi(x) \wedge \neg\psi(x)$ is satisfiable in $(A, a)_{a \in C \cup I}$.*

Let C_0 be any countable subset of C. For each $L(C_0 \cup I)$ formula $\varphi(x)$ that is satisfiable in $(A, a)_{a \in C_0 \cup I}$ let ψ_φ be one of the formulas satisfying (*) for φ. Since $C_0 \cup I$ is countable, there is a countable set C_1 such that $C_0 \subseteq C_1 \subseteq C$ and such that the parameters of ψ_φ are in C_1. Continuing in this way to create sets C_k, let $C' = \bigcup\{C_k : k \in \mathbb{N}\}$. Let $p'(x)$ be restriction of $p(x)$ to C'. We have (**): *for each $L(C' \cup I)$-formula $\varphi(x)$ that is satisfiable in $(A, a)_{a \in C' \cup I}$ there exists $\psi(x) \in p'(x)$ such that $\varphi(x) \wedge \neg\psi(x)$ is satisfiable in $(A, a)_{a \in C' \cup I}$.* Note also that $(a_k : k \in \mathbb{N})$ is a sequence of indiscernibles in $(A, a)_{a \in C'}$. □

A theorem implying Morley's theorem, proof continued

Proof.

By the Standard Lemma there is a model B of $\mathrm{Th}((A, a)_{a \in C'})$ that contains a family $(b_\alpha : \alpha < \kappa)$ realising the Ehrenfeucht-Mostowski type of $(a_k : k \in \mathbb{N})$. We may assume this model is of the form $(B, a)_{a \in C'}$. Using the Third Goal we know that there is an elementary substructure B' of B which is atomic over $C' \cup \{b_\alpha : \alpha < \kappa\}$.

The proof will be finished once we show that $p'(x)$ is not realised in $(B', a)_{a \in C'}$. For then the downward Löwenheim-Skolem Theorem implies that B' has an elementary substructure B'' of cardinality κ which contains C'. Then B'' is a model of cardinality κ which is not κ-saturated. (In fact, it is not even ω_1-saturated.) □

A theorem implying Morley's theorem, proof finished

Claim

The type $p'(x)$ is not realised in $(B', a)_{a \in C'}$.

Proof.

Recall that we have (**): *for each $L(C' \cup I)$-formula $\varphi(x)$ that is satisfiable in $(A, a)_{a \in C' \cup I}$ there exists $\psi(x) \in p'(x)$ such that $\varphi(x) \wedge \neg\psi(x)$ is satisfiable in $(A, a)_{a \in C' \cup I}$.*

So suppose $p'(x)$ is realised in $(B', a)_{a \in C'}$ by some tuple b. We have that $\mathrm{tp}_{B'}(b / C' \cup \{b_\alpha : \alpha < \kappa\})$ is isolated so it contains a complete formula $\varphi(x, b_{\alpha_0}, \ldots, b_{\alpha_n})$. So we have that $\varphi(x, b_{\alpha_0}, \ldots, b_{\alpha_n}) \rightarrow \psi(x)$ holds in B' for every $\psi(x) \in p'(x)$. But since $b_{\alpha_0}, \ldots, b_{\alpha_n}$ and $a_0, \ldots, a_n$ realize the same Ehrenfeucht-Mostowski type over C', we have that $\varphi(x, a_0, \ldots, a_n) \rightarrow \psi(x)$ is valid in A for each formula $\psi(x) \in p'(x)$. But that contradicts (**). □

Morley's Theorem

Morley's Theorem

If a countable theory T is λ-categorical for an uncountable cardinal λ, then it is λ-categorical for all uncountable cardinal λ.

End of the course. And Merry Christmas and Happy New Year!