Section 3

Compactness Theorem

Compactness Theorem

The most important result in model theory is:

Compactness Theorem

Let T be a theory in language L. If every finite subset of T has a model, then T has a model.

It follows from the completeness theorem for first-order logic, but we can also prove it by purely model-theoretic means.

For convenience call a theory T *finitely consistent* if any finite subset of T has a model. We have to show that finitely consistent implies consistent.

Proof of Compactness Theorem

Lemma

Suppose T is an L-theory with the following properties:

1. T is finitely consistent.
2. For any formula φ the theory T contains either φ or $\neg\varphi$.
3. If T contains a sentence $\exists x\, \varphi(x)$, then there is a closed term t such that T also contains $\varphi(t)$.

Then T has a model.

Proof.

Note that if T_0 is a finite subset of T and $T_0 \models \varphi$, then $\varphi \in T$ by the first two properties. We construct a model M by taking the closed terms in the language in L and identifying s and t whenever $s = t \in T$ (this is an equivalence relation), and we say that $([t_1], \ldots, [t_n]) \in R^M$ in case $R(t_1, \ldots, t_n) \in T$ and $f^M([t_1], \ldots, [t_n]) = [f(t_1, \ldots, t_n)]$ (this is well-defined). Now show by induction on the structure of the term t that $t^M = [t]$ and the structure of the formula φ that $M \models \varphi$ if and only if $\varphi \in T$.

Proof of Compactness Theorem, continued

Lemma

Suppose T is finitely consistent. Then T can be extended to a theory T' which is finitely consistent and which for any sentence φ contains either φ or $\neg\varphi$.

Proof.

Use Zorn's Lemma to extend T to a maximal finitely consistent theory T'. We will show that if such a T' does not contain a formula φ, it has to contain $\neg\varphi$. So suppose $\varphi \notin T'$. Note that because T' is maximal, the theory $T' \cup \{\varphi\}$ cannot be finitely consistent, meaning that there is a finite subtheory $T_0 \subseteq T'$ such that $T_0 \cup \{\varphi\}$ has no models.

It now follows that $T' \cup \{\neg\varphi\}$ is finitely consistent. For let $T_1 \subseteq T'$ be a finite. Since $T_0 \cup T_1 \subseteq T'$ is finite and T' is finitely consistent, $T_0 \cup T_1$ has a model M. As M models T_0 it cannot model φ and must model $\neg\varphi$. So $T_1 \cup \{\neg\varphi\}$ is consistent. But if $T' \cup \{\neg\varphi\}$ is finitely consistent and T' is maximal, then $T' \cup \{\neg\varphi\} = T'$. So $\neg\varphi \in T'$, as desired. □

Proof of Compactness Theorem, finished

Lemma

Suppose T is a finitely consistent L-theory. Then L can be extended to a language L' and T to a finitely consistent L'-theory T' such that for any L-sentence of the form $\exists x\, \varphi(x)$ in T' there is a term t in L' such that $\varphi(t) \in T'$.

Proof.

For any sentence of the form $\exists x\, \varphi(x)$ which belongs to T we add a fresh constant c to the language L' and the sentence $\varphi(c)$ to the theory T'. □

Proof of the Compactness Theorem

Starting from any finitely consistent theory T we can, by alternatingly applying the previous two lemmas, create a theory which has all the three properties of the first lemma. So T has a model.

Section 4

The method of diagrams

Diagrams

Definition

If M is a model in a language L, then the collection of quantifier-free L_M-sentences 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 $T_\forall = \{ \varphi : T \models \varphi$ and φ is universal $\}$. Clearly, $T \models T_\forall$. We need to prove the converse.

So suppose M is a model of $T_\forall$. It suffices 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 $M \models T_\forall$ where $T_\forall = \{\, \varphi : T \models \varphi$ 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 sentences $\psi_1, \ldots, \psi_n \in \mathrm{Diag}(M)$ which are inconsistent with T. Write $\psi = \bigwedge_i \psi$ and note that $\psi \in \mathrm{Diag}(M)$ and ψ is already inconsistent with T. Replace the constants from M in ψ by variables $x_1, \ldots, x_n$ and we obtain ψ'; because the constants from M do not appear in T, the theory T is already inconsistent with $\exists x_1, \ldots, x_n \, \psi'$. So $T \models \neg\exists x_1, \ldots, x_n \, \psi'$ and $T \models \forall x_1, \ldots, x_n \, \neg\psi'$. Therefore the sentence $\forall x_1, \ldots, x_n \, \neg\psi'$ belongs to $T_\forall$; but it is false in M as M is a model of ψ. This contradicts the assumption that $M \models T_\forall$. □

Chang-Łoś-Suszko Theorem

Definition

A $\forall\exists$-*sentence* is a sentence which consists of a sequence of universal quantifiers, then a sequence of existential quantifiers and then a quantifier-free formula. A theory T can be *axiomatised by* $\forall\exists$-*sentences* if there is a set T' of $\forall\exists$-sentences which has the same models as T.

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. T is *preserved by unions of chains* if for any chain of models of T and embeddings between them, also the colimit is a model of T.

Chang-Łoś-Suszko Theorem

The following statements are equivalent:

(1) T is preserved by directed unions.

(2) T is preserved by unions of chains.

(3) T can be axiomatised by $\forall\exists$-sentences.

Chang-Łoś-Suszko Theorem, proof

Proof. We just show (2) $\Rightarrow$ (3). Suppose T is preserved by unions of chains. Again, let

$$T_{\forall\exists} = \{\varphi : \varphi \text{ is a } \forall\exists\text{-sentence and } T \models \varphi\},$$

and let B be a model of $T_{\forall\exists}$. We will construct a chain of embeddings

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

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 $T_{\forall\exists}$.

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. □

Section 5

Theorems of Robinson, Craig and Beth

Robinson's Consistency Theorem

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 $T_1 \cup T_2$.

First lemma

Lemma

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

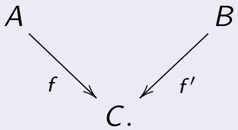

Proof. Consider $T = \mathrm{ElDiag}(A) \cup \mathrm{ElDiag}(B)$ (making sure we use different constants for the elements from A and B!). We need to show T has a model; so suppose T is inconsistent. Then, by Compactness, a finite subset of T has no model; taking conjunctions, we have sentences $\varphi(a_1, \ldots, a_n) \in \mathrm{ElDiag}(A)$ 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 $A \equiv B \restriction L$. $\square$

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

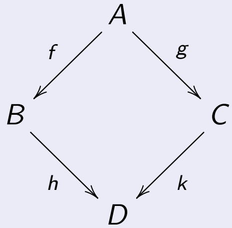

where D is an L'-structure, h is an L-elementary embedding and k is an L'-elementary embedding.

Proof.

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

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

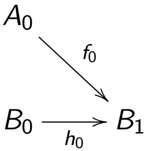

with h_0 an L_2-elementary embedding and f_0 an L-elementary embedding. Now by applying the second lemma to f_0 and the identity on A, we obtain

...

Robinson's consistency theorem, proof finished

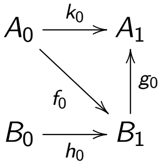

where g_0 is L-elementary and k_0 is L_1-elementary. Continuing in this way we obtain a diagram

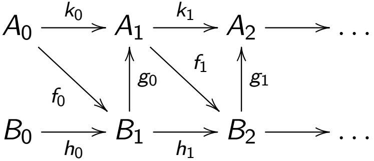

where the k_i are L_1-elementary, the f_i and g_i are L-elementary and the h_i are L_2-elementary. The colimit C of this directed system is both the colimit of the A_i and of the B_i. So A_0 and B_0 embed elementarily into C by the elementary systems lemma; hence C is a model of both T_1 and T_2, as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

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

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

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

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

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

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

(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

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

So there is a Craig Interpolant θ such that

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

By symmetry also

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

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