Linstrom's Theorem Lecture

by Nathanael Leedom Ackerman

October 2, 2006

TALK SLOWLY AND WRITE NEATLY!!

0.1 Introduction

0.1.1 Introduction

In this talk I will prove Lindstrom's Theorem which says

Linstrom's Theorem

Theorem 0.1.1.1 (Linstrom's Theorem)**.** *First order logic ($\mathcal{L}_{\omega,\omega}$) is a maximal logic among those logics which satisfy the Lowenheim-Skolem property for countable sets of sentences, compactness and have basic closure properties.*

Now before I can prove this I have to define what these things mean.

0.1.2 Arbitrary Logics

Vocabulary

Definition 0.1.2.1. We say that τ is a Vocabulary if $\tau = \langle \boldsymbol{S}_\tau, \boldsymbol{R}_\tau, \boldsymbol{F}_\tau \rangle$ where

$\boldsymbol{S}_\tau = \{s_i : i \in I\}$ where each s_i is a sorts.

$\boldsymbol{R}_\tau = \{R_j : j \in J\}$ where each each R_j is a relation symbol with a finite number of arguments such that each argument is equipped with a sort $s \in \boldsymbol{S}$

$\boldsymbol{F}_\tau = \{f_k : k \in K\}$ where each each f_k is a function symbol with a finite number of arguments such that finitary function symbol where each argument and the function symbol itself is equipped with a sort $s \in \boldsymbol{S}$

Renaming

Definition 0.1.2.2. Let τ, σ be vocabularies. We say $\rho : \tau \to \sigma$ is a Renaming if ρ is a bijection which takes sorts to sort, functions to functions, relations to relations and which is consistent with the sorts assigned to relation and function symbols. (i.e. if $sort(f) = s$ then $sort(\rho(f)) = \rho(s)$)

Structures of a Vocabulary

Definition 0.1.2.3. Let τ be a vocabulary. We say that $\mathcal{A}$ is a τ-structure ($\mathcal{A} \in Str[\tau]$) if

(1) For each $s \in \boldsymbol{S}_\tau$ there is a set $A_s \in \mathcal{A}$

(2) For each $R \in \boldsymbol{R}_\tau$ of sorts $(s_1, \cdots, s_n)$ there is a set $R^{\mathcal{A}} \in \mathcal{A}$ such that $R^{\mathcal{A}} \subseteq A_{s_1} \times \cdots A_{s_n}$

(3) For each $f \in \boldsymbol{F}_\tau$ of sorts $(s_0, s_1, \cdots, s_n)$ there is a function $f^{\mathcal{A}} \in \mathcal{A}$ such that $f^{\mathcal{A}} : A_{s_1} \times \cdots A_{s_n} \to A_{s_0}$

In this case we say that $\tau = \tau_{\mathcal{A}}$ (so we always have $\mathcal{A} \in Str[\tau_{\mathcal{A}}]$)

Reduct

Definition 0.1.2.4. Let $\sigma \subseteq \tau$ be vocabularies and let $\mathcal{A} \in Str[\tau]$ be a τ structure. Then we define the reduct of $\mathcal{A}$ to σ $(\mathcal{A}|\sigma)$ to be gotten by simply ignoring all $f, R, s \in \tau - \sigma$ in the structure $\mathcal{A}$

We are now ready to define an abstract logic.

Abstract Logic

Definition 0.1.2.5. An abstract logic is a pair $(L, \models_L)$ such that

L is a mapping from vocabularies (so that $L[\tau]$ is the class of L-sentences of vocabulary τ) such that

If $\sigma \subseteq \tau$ are vocabularies then $L[\sigma] \subseteq L[\tau]$

$\models_L$ (the L-satisfaction relation) is a relation between structures and L-sentences such that

(1) If $\mathcal{A} \models_L \varphi$ then $\varphi \in L[\tau_{\mathcal{A}}]$

(2) (Isomorphism Property) If $\mathcal{A} \models_L \varphi$ and $\mathcal{A} \cong \mathcal{B}$ then $\mathcal{B} \models_L \varphi$

(3) (Reduct Property) If $\varphi \in L[\tau]$ and $\tau \subseteq \tau_{\mathcal{A}}$ then

$$\mathcal{A} \models_L \varphi \text{ if and only if } \mathcal{A}|\tau \models_L \varphi$$

(4) (Renaming Property) Let $\rho : \tau \longrightarrow \sigma$ be a renaming. Then for each $\varphi \in L[\tau]$ there is a $\varphi^\rho \in L[\sigma]$ such that for each $\mathcal{A} \in Str[\tau]$

$$\mathcal{A} \models_L \varphi \text{ if and only if} \mathcal{A}^\rho \models \varphi^\rho$$

where $\mathcal{A}^\rho$ is the image under ρ of the structure $\mathcal{A}$ (and so $\mathcal{A}^\rho \in Str[\sigma]$)

Equivalence in an Abstract Logic

Definition 0.1.2.6. Let $(L, \models_L)$ be an abstract logic and let $A, B \in Str[\tau]$. Then we say A, B are L-equivalent ($A \equiv_L B$) if

$$(\forall \varphi \in L[\tau]) A \models_L \varphi \text{ if and only if } B \models_L \varphi$$

Notice that in this definition of abstract logic the only information which is important is which models satisfy which formulas. Specifically we can think about an abstract formula as simply the collection of models which satisfy it. And modulo some information concerning vocabularies, an abstract logic is just a collection of such formulas.

Formulas as collections of models

Definition 0.1.2.7. Given a logic $(L, \models_L)$ we define $Mod_L^\tau(\varphi) = \{M \in Str[\tau] : M \models_L \varphi\}$. We call any such $Mod_L^\tau(\varphi)$ and elementary class of L ($Mod_L^\tau(\varphi) \in EC(L)$).

This way of viewing abstract logics gives rise to a natural partial on the class of all such logics.

Comparing Strengths of Abstract Logics

Definition 0.1.2.8. Let $(L, \models_L), (L^*, \models_{L^*})$ be abstract logics. We then say $L \leq L^*$ if

$$(\forall \tau)(\forall \varphi \in L[\tau])(\exists \varphi^* \in L^*[\tau]) Mod_L^\tau(\varphi) = Mod_{L^*}^\tau(\varphi^*)$$

Intuitively we say $L \leq L^*$ if for every formula of L there is a formula of L^* equivalent to it. (And so L^* has at least as much expressive power as L).

0.1.3 Other Definitions

Basic Closure Properties

Definition 0.1.3.1. Let $(L, \models_L)$ be a logic. We say that L has Basic Closure Properties if

(i) $(\forall \varphi, \psi \in L[\tau])(\exists \varphi \wedge \psi \in L[\tau])$ such that

$$Mod_L^\tau(\varphi) \cap Mod_L^\tau(\psi) = Mod_L^\tau(\varphi \wedge \psi)$$

(L is closed under conjunction)

(ii) $(\forall \varphi \in L[\tau])(\exists \neg \varphi \in L[\tau])$ such that

$$Mod_L^\tau(\varphi) = \{M : M \in Str[\tau] \wedge M \notin Mod_L^\tau(\varphi \wedge \psi)\}$$

(L is closed under negation)

These are necessary in order to allow us to do basic manipulations of formulas.

Lowenheim-Skolem property

Definition 0.1.3.2. We say that a logic $(L, \models_L)$ has the Lowenheim-Skolem property for countable collections of sentences if for all countable collections of formulas $\Phi \subseteq L[\tau]$ we have

$$(\exists M \in Str[\tau])(\forall \varphi \in \Phi) M \models_L \varphi \Rightarrow (\exists N \in Str[\tau])(\forall \varphi \in \Phi) N \models \varphi \wedge |N$$

What this is says is that if a countable collection of sentences has a model, then it must have a countable model .

Compactness

Definition 0.1.3.3. We say that a logic $(L, \models_L)$ is compact if for all vocabularies τ

$$(\forall S \subseteq L[\tau]) \bigcap_{\varphi \in S} Mod_L^\tau(\varphi) = \emptyset \Rightarrow (\exists S' \subseteq S)(\bigcap_{\varphi \in S'} Mod_L^\tau(\varphi) = \emptyset) \wedge |S'| <$$

In other words, a logic is compact if given any collection of sentences, that sentence has a model if and only if every finite subset of those sentences has a model.

This is in fact one of the most important properties of fist order logic.

0.2 Proof

Linstrom's Theorem

Theorem 0.2.0.4 (Linstrom's Theorem)**.** *Let* $(L, \models_L)$ *be a logic such that*

(a) L satisfies the basic closure properties.

(b) L satisfies the Lowenheim-Skolem property for countably many formulas.

(c) L satisfies compactness.

(d) $L \geq \mathcal{L}_{\omega,\omega}$ *(first order logic)*

Then $L \equiv \mathcal{L}_{\omega,\omega}$

Proof. We are now ready to prove Linstrom's theorem. The proof will proceed in four steps.

Compactness/LS-Property is preserved by sublo

Claim 0.2.0.5. *Let $(L, \models)$ be a logic with compactness and the Lowenheim-Skolem property. If $L' \subseteq L$ then L' has compactness and the Lowenheim-Skolem property as well.*

Proof. Compactness: Lets assume that L satisfies compactness and $L' \subseteq L$.

Now lets let $S \subseteq L'$ be such that every finite subset of S has a model. Then we know that $S \subseteq L$ and hence by compactness of L S has a model. So, because S was arbitrary we see that L' is compact.

Lowenheim-Skolem Property: Lets assume that L satisfies the Loenheim-Skolem property for countable sentences and $L' \subseteq L$.

Now lets let $S \subseteq L'$ be a countable collection of sentences with a model. Then we know that $S \subseteq L$ and hence by the Lowenheim-Skolem property of L S has a countable model. So, because S was arbitrary we see that L' has the Lowenheim-Skolem property. □

Smallest Logic with Basic Closure Properties

Claim 0.2.0.6. *Let M be an abstract logic. Then there is a smallest abstract logic $\overline{M}$ satisfying the following:*

(i) $M \subseteq \overline{M}$

(ii) $\overline{M}$ has basic closure properties

And further if N is a collection of formulas satisfying (i) and (ii) then $\overline{M} \subseteq N$

Proof. $M[\sigma]$ is just a collection of subsets of all models with signature σ. So if we let $\overline{M}[\sigma] = \{\varphi : \varphi$ can be gotten from sets in $M[\sigma]$ by a finite number of intersections and complementations$\}$ then it is easy to check that $\overline{M}$ satisfies the above properties. $\square$

Now to get a contradiction lets assume that there is a $\psi \in L[\tau] - \mathcal{L}_{\omega,\omega}[\tau]$. By the previous two claims we know that $\overline{\mathcal{L}_{\omega,\omega} \cup \{psi\}}$ satisfies

- Compactness
- The Lowenheim-Skolem property
- Basic Closure Properties
- $\mathcal{L}_{\omega,\omega} \leq \overline{\mathcal{L}_{\omega,\omega} \cup \{psi\}}$

And in particular satisfies the conditions for Lindstrom's theorem. So with out loss of generality we can assume that $L = \overline{\mathcal{L}_{\omega,\omega} \cup \{psi\}}$ **Finite Vocabulary**

Claim 0.2.0.7. *Let $\varphi \in L[\sigma]$. Then there is a finite*

$\sigma_0 \subseteq \sigma$ *such that for any* σ *structures* A, B

$A|\sigma_0 \cong B|\sigma_0$ *implies* ($A \models_L \psi$ *if and only if* $B \models_L \varphi$)

Proof. Let σ' be a disjoint copy of σ and let $\Phi \subseteq L[\sigma \cup \sigma']$ be the set

$$\begin{aligned} \Psi &= \{(\forall x_1, \ldots, x_n)(R(x_1, \ldots, x_n) \leftrightarrow R'(x_1, \ldots, x_n)) : R \in \sigma, n - ary \\ &\cup \{(\forall x_1, \ldots, x_n)(f(x_1, \ldots, x_n) \leftrightarrow f'(x_1, \ldots, x_n)) : f \in \sigma, n - ary\} \end{aligned}$$

Now clearly $\Phi \models \psi \leftrightarrow \psi'$ (by renaming). But that means that

$$\bigcap \Phi \cap \{\neg(\psi \leftrightarrow \psi')\} = \emptyset$$

So by compactness there must be a finite set $\Phi_0 \subseteq \Phi$ such that

$$\bigcap \Phi_0 \cap \{\neg(\psi \leftrightarrow \psi')\} = \emptyset$$

and hence such that $\Phi_0 \models \psi$.

This is just another form of compactness which says, essentially, that if you can prove a sentence from a collection of sentences then you have to be able to prove it from a finite subcollection.

But we then also know that $\Phi \subseteq L[\sigma_0 \cup \sigma_0']$ for some finite σ_0 (because Φ is finite).

Now suppose we have A, B such that $A|\sigma_0 \cong B|\sigma_0$. We can then define a model C of signature $\sigma \cup \sigma'$ such that $C|\sigma \cong A$ and $C|\sigma' \cong B$. By our assumption we know that $C \models \Phi_0$ and hence we know that $A \cong C|\sigma \models \psi$ if and only if $B \cong C|\sigma' \models \psi'$. $\square$

This simply says that if we want to know about any particular formula of L we only have to look at a finite vocabulary.

So in particular we can choose a τ_0 according to the previous theorem with respect to $\psi \in L[\tau] - \mathcal{L}_{\omega,\omega}[\tau]$

Claim 1

Claim 0.2.0.8. *There are τ-structures $\mathcal{A}, \mathcal{B}$ such that*

$$\mathcal{A} \models_L \psi \text{ and } \mathcal{B} \models_L \neg\psi$$

$$\mathcal{A}|\tau_0 \equiv_L \mathcal{B}|\tau_0$$

$$A = B$$

Proof. Let $\varphi_1, \varphi_2, \ldots,$ be a complete list of the $\mathcal{L}_{\omega,\omega}[\tau_0]$ sentences. By induction we obtain a sequence $\psi_1, \psi_2, \ldots$ such that for each n

- $\psi_n = \varphi_n$ or $\psi_n = \neg\varphi_n$
- $\psi \wedge \psi_1 \wedge \cdots \wedge \psi_n$ is not equivalent to a first order sentence.

We then also have that $\neg\psi \wedge \psi_1 \wedge \cdots \wedge \psi_n$ also must not be equivalent to a first order sentence (because if it equivalent to a first order sentence η then ψ would be equivalent to $\neg\eta \wedge \neg\psi_1 \wedge \cdots \wedge \neg\psi_n$)

So in particular both $\psi \wedge \psi_1 \wedge \cdots \wedge \psi_n$ and $\neg\psi \wedge \psi_1 \wedge \cdots \wedge \psi_n$ are satisfiable (i.e. not equivalent to $R \wedge \neg R$).

Let $\Phi = \{\psi_n : n \in \omega\}$. Hence by compactness there are models $\mathcal{A} \models \Phi \cup \{\psi\}$ and $\mathcal{B} \models \Phi \cup \{\neg\psi\}$ and by the Lowenheim-Skolem property we can assume that $\mathcal{A}$ and $\mathcal{B}$ are countable. (And they have the same first order theory Φ).

It is also simply to assume that the domain A of $\mathcal{A}$ is the same set as the domain B of $\mathcal{B}$ (i.e. $A = B$). $\square$

Claim 2

Claim 0.2.0.9. *There are τ-structures A, B such that*

$$\mathcal{A} \models_L \psi \text{ and } \mathcal{B} \models_L \neg\psi$$

$$\mathcal{A}|\tau_0 \cong \mathcal{B}|\tau_0$$

$A = B$

Proof. First we will need to consider a vocabulary $\tau^* = \tau \cup \tau' \cup \{f_n, g_n : n \in \omega\}$ (where τ' is a copy of τ and f_n, g_n are $2n+1$-ary function symbols.

Next for each $n \in \omega$ fix an enumeration $\langle \xi(x_1, \ldots, x_n, x) : i \in \omega \rangle$ f all $\mathcal{L}_{\omega,\omega}[\tau_0]$-formula with free variables among $x_1, \ldots, x_n, x$.

Let Γ consist of

- $\psi, \neg\psi'$ (i.e. the τ reduct is a model of ψ and the τ' reduct is a model of $\neg\psi'$)
- $\varphi \leftrightarrow \varphi'$ for each sentence $\varphi \in \mathcal{L}_{\omega,\omega}[\tau_0]$ (i.e. The τ_0 reduct and the τ_0' reduct are elementarily equivalent)

$(*_{n,r})$ $(\forall\overline{x}\forall\overline{y}\forall x)\,(\exists y \bigwedge_{i=0}^{r}(\xi_i(\mathbf{x}, x) \leftrightarrow \xi_i'(\mathbf{y}, y))) \rightarrow (\bigwedge_{i=0}^{r}(\xi_i(\mathbf{x}, x) \leftrightarrow \xi_i'(\mathbf{y},$

$(*_{n,r})$ $(\forall\overline{x}\forall\overline{y}\forall y)\,(\exists x \bigwedge_{i=0}^{r}(\xi_i(\mathbf{x}, x) \leftrightarrow \xi_i'(\mathbf{y}, y))) \rightarrow (\bigwedge_{i=0}^{r}(\xi_i(\mathbf{x}, g_n(\mathbf{x}, \mathbf{y}, y))$

where the in the last two conditions $n, r \in \omega$ and $\mathbf{x} = x_1, \ldots, x_n$ and $\mathbf{y} = y_1, \ldots, y_n$.

The intended purpose of f_n and g_n is to allow us to create an isomorphism from the τ-reduct of our model to the τ'-reduct of our model to τ'.

Claim 0.2.0.10. Γ *has a model*

Proof. Suppose not. Then there must be a finite collection of sentences $\Gamma_0 \subseteq \Gamma$ which doesn't have a model (by compactness). But, if Γ_0 is finite then it contains only a finite number of $(*_{n,r})$. But we know that for each n $\models (*_{n,r+1}) \rightarrow (*_{n,r})$. So we have that $\Gamma_0 \leftrightarrow \Phi_0$ where

- $\Phi_0 \subseteq \Gamma$
- For each f_n, g_n there is at most a single $(*_{n,r}) \in \Phi_0$

But, we know by the previous claim that every axiom except for the $(*)$ has a model. And, because each of $(*_{n,r})$ is a Skolem sentence, and because for each function there is at most one of them in Φ_0, our model from the previous claim can be extended to be a model for Φ_0 (Note that a Skolem sentence is of the form $(\exists x)\varphi(x, \mathbf{y}) \to \varphi(f_\varphi(\mathbf{y}), y)$. The value of f_φ is then found by the axiom of choice).

But we have that $\Phi_0 \leftrightarrow \Gamma_0$ and hence Γ_0 has a model. $\Rightarrow\Leftarrow$ we choose Γ_0 so that it wouldn't have a model.

Hence Γ has a model. $\square$

Notice though that this proof does require the countable axiom of choice.

Now lets let $\mathcal{D}$ be a model of Γ and lets let $\mathcal{A} = \mathcal{D}|\tau$

and $\mathcal{B} = (\mathcal{D}|\tau')^{-'}$ where $' : \tau \to \tau'$ is the given renaming of τ.

We then clearly have

- $A = B = D$
- $\mathcal{A} \models \psi$
- $\mathcal{B} \models \neg\psi$
- $\mathcal{A}|\tau_0 \equiv \mathcal{B}|\tau_0$

So if we can show that $\mathcal{A}|\tau_0 \cong \mathcal{B}|\tau_0$ then we are done.

Let $d_1, d_2, \ldots$ be an enumeration of D. Since $\mathcal{A}|\tau_0 \equiv \mathcal{B}|\tau_0$ we have by $(*_{n,r})$

$$\begin{aligned}(\mathcal{A}|\tau_0, d_1) &\equiv (\mathcal{B}|\tau_0, f_0(d_1)) \\ (\mathcal{A}|\tau_0, d_1, g_1(d_1, f_0(d_1), d_1)) &\equiv (\mathcal{B}|\tau_0, f_0(d_1), d_1) \\ (\mathcal{A}|\tau_0, d_1, g_1(d_1, f_0(d_1), d_1), d_2) &\equiv (\mathcal{B}|\tau_0, f_0(d_1), d_1, f_1(\ldots), \ldots))\end{aligned}$$

Continuing in this way we see that we can obtain sequences $a_1, a_2, \ldots, b_1, b_2, \ldots$ such that

- $A = \{a_n : n \in \omega\}$
- $B = \{b_n : n \in \omega\}$
- For each n, $(\mathcal{A}|\tau_0, a_1, \ldots, a_n) \equiv (\mathcal{B}|\tau_0, b_1, \ldots, b_n)$

But then if we define $\pi : \mathcal{A}|\tau_0 \rightarrow \mathcal{B}|\tau_0$ we see

- π is a bijection
- $\mathcal{A}|\tau_0 \models \varphi(\overline{a})$ if and only if $\mathcal{B}|\tau_0 \models \varphi(\overline{b})$ for any quantifier free formula of τ_0.

Hence, π is an isomorphism. And so we are done with the claim. $\square$

But this is a contradiction and hence $\psi \in \mathcal{L}_{\omega,\omega}$ and so $L \equiv \mathcal{L}_{\omega,\omega}$. $\square$