Strong conceptual completeness and internal adjunctions in **Def**(T)

Jesse Han

McMaster University

October 10, 2017

What is strong conceptual completeness for first-order logic?

- A strong conceptual completeness statement for a logic(al doctrine) is an assertion that a theory in this logic(al doctrine) can be recovered from an appropriate structure formed by the models of the theory.
- Makkai proved such a theorem for first-order logic showing one could reconstruct a first-order theory T from $\mathbf{Mod}(T)$ equipped with structure induced by taking ultraproducts.
- Before we dive in, let's look at a well-known theorem from model theory, with the same flavor, which Makkai's result generalizes: the Beth definability theorem.
- (This theorem says something like: if there is only a unique way to expand the models of a theory by a new predicate, then this predicate must have been definable in the theory to begin with.)

The Beth theorem

Theorem.

Let $L_0 \subseteq L_1$ be an inclusion of languages with no new sorts. Let T_1 be an L_1-theory. Let $F : \mathbf{Mod}(T_1) \to \mathbf{Mod}(\varnothing_{L_0})$ be the reduct functor. Suppose you know any of the following:

1. *There is a L_0-theory T_0 and a factorization:*

$$\begin{array}{ccc} \mathbf{Mod}(T_1) & \xrightarrow{F} & \mathbf{Mod}(\varnothing_{L_0}) \\ & \underset{\simeq}{\searrow} & \uparrow \\ & & \mathbf{Mod}(T_0) \end{array}$$

2. *F is full and faithful.*
3. *F is injective on objects.*
4. *F is full and faithful on automorphism groups.*
5. *F is full and faithful on $\mathrm{Hom}_{L_1}(M, M^{\mathcal{U}})$ for all $M \in \mathbf{Mod}(T_1)$ and all ultrafilters $\mathcal{U}$.*
6. *Every L_0-elementary map is an L_1-homomorphism of structures.*

<u>Then:</u> (!) Every L_1-formula is T_1-provably equivalent to an L_0-formula.

Useful consequence of Beth's theorem

Corollary.

Let T be an L-theory, let $\overline{S}$ be a finite product of sorts. Let $X : \mathbf{Mod}(T) \to \mathbf{Set}$ be a subfunctor of $M \mapsto \overline{S}(M)$.

Then: if X commutes with ultraproducts on the nose ("satisfies a Łos' theorem"), then X was definable, i.e. X is an evaluation functor for some definable set $\varphi \in \mathbf{Def}(T)$.

Proof.

(Sketch): expand each model M of T by a new sort $X(M)$. Use commutation with ultraproducts to verify this is an elementary class. Then we are in the situation of $1 \implies (!)$ from Beth's theorem. □

How does strong conceptual completeness enter this picture?

- Plain old conceptual completeness (this was one of the key results of Makkai-Reyes) says that if an interpretation $I : T_1 \to T_2$ induces an equivalence of categories $\mathbf{Mod}(T_1) \overset{I^*}{\simeq} \mathbf{Mod}(T_2)$, then I must have been a bi-interpretation.
  So, it proves $1 \implies (!)$, and therefore the corollary.
- Strong conceptual completeness is the following upgrade of the corollary.

Strong conceptual completeness, I

Theorem.

Let T be an L-theory. Let X be any functor $\mathbf{Mod}(T) \to \mathbf{Set}$. Suppose that you have:

- *for every ultraproduct $\prod_{i \to \mathcal{U}} M_i$ a way to identify $X(\prod_{i \to \mathcal{U}} M_i) \overset{\Phi_{(M_i)}}{\simeq} \prod_{i \to \mathcal{U}} X(M_i)$ ("there exists a transition isomorphism"), such that*
- *(X, Φ) preserves ultraproducts of models/elementary embeddings ("is a pre-ultrafunctor"), and also*
- *preserves all canonical maps between ultraproducts ("preserves ultramorphisms").*

Then: there exists a $\varphi(x) \in T^{\mathrm{eq}}$ such that $X \simeq \mathrm{ev}_{\varphi(x)}$ as functors $\mathbf{Mod}(T) \to \mathbf{Set}$. (We call such X an ultrafunctor.)

Strong conceptual completeness, II

Now, what's changed between this statement and that of the useful corollary to Beth's theorem?

- We dropped the *subfunctor* assumption! We don't have such a nice way of knowing exactly how $X(M)$ is obtained from M. We only have the invariance under ultra-stuff. We've left the placental warmth of the ambient models and we're considering some kind of abstract permutation representation of $\mathbf{Mod}(T)$.
- Yet, if X respects enough of the structure induced by the ultra-stuff, then X must have been constructible from our models in some first-order way ("is definable").
- (With this new language, the corollary becomes: "strict sub-pre-ultrafunctors of definable functors are definable.")

Strong conceptual completeness, III

Actually, Makkai proved something more, by doing the following:

- Introduce the notions of ultracategory and ultrafunctors by requiring all this extra ultra-stuff to be preserved.
- Develop a general duality theory between pretoposes ("$\mathbf{Def}(T)$") and ultracategories ("$\mathbf{Mod}(T)$") via a contravariant 2-adjunction ("generalized Stone duality").
- In particular, from this adjunction we get $\mathbf{Pretop}(T_1, T_2) \simeq \mathbf{Ult}(\mathbf{Mod}(T_2), \mathbf{Mod}(T_1))$.

Therefore, SCC tells us how to recognize a reduct functor in the wild between two categories of models—i.e., if there is some uniformity underlying a functor $\mathbf{Mod}(T_2) \to \mathbf{Mod}(T_1)$ due to a purely syntactic assignment $T_1 \to T_2$. Just check if the ultra-structure is preserved!

Caveat. Of course, one has an infinite list of conditions to verify here.

- So the only way to actually do this is to recognize some kind of uniformity in the putative reduct functor which lets you take care of all the ultramorphisms at once.
- But it gives you another way to think about uniformities you need.
- It also gives you a way to check that something can never arise from any interpretation!

Ultramorphisms, I

- Part of the critera for (X, Φ) (a functor $X : \mathbf{Mod}(T) \to \mathbf{Set}$ plus a choice of transition isomorphism Φ) to be definable was "preserving ultramorphisms."
- What are ultramorphisms? Loosely speaking, ultraproducts are a kind of universal construction in **Set**, and so there are certain canonical comparison maps between them induced by their universal properties. (By the Los theorem, these things are "absolute" in the sense that no matter what first-order structure you put on a set, these maps will always be elementary embeddings.)
- Out of mercy, I will spare you the formal definition (because then I'd have to define ultragraphs, ultradiagrams, and ultratransformations...)
- Keep in mind these two examples:

Ultramorphisms, II

Examples.

- *The diagonal embedding into an ultrapower.*
- *Generalized diagonal embeddings. More generally, let $f : I \to J$ be a function, let $\mathcal{U}$ be an ultrafilter on I and let $\mathcal{V}$ be the pushforward ultrafilter on J. Then for any I-indexed sequence of structures $(M_i)_{i \in I}$, there is a canonical map $\delta_f : \prod_{j \to \mathcal{V}} M_{f(i)} \to \prod_{i \to \mathcal{U}} M_i$ given by taking the diagonal embedding along each fiber of f.*

Δ-functors induce continuous maps on automorphism groups

- Why should we expect ultramorphisms to help us identify evaluation functors in the wild?
- Here's an result which might indicate that knowing that they're preserved tells us something nontrivial.

Definition.

Say that $X : \mathbf{Mod}(T) \to \mathbf{Mod}(T')$ is a Δ-functor if it preserves ultraproducts and diagonal maps into ultrapowers.

Equip automorphism groups with the topology of pointwise convergence.

Theorem.

If X is a Δ-functor from $\mathbf{Mod}(T)$ to $\mathbf{Mod}(T')$, then X restricts to a continuous map $\mathrm{Aut}(M) \to \mathrm{Aut}(X(M))$ for every $M \in \mathbf{Mod}(T)$.

Proof.

- The topology of pointwise convergence is sequential, so to check continuity it suffices to check convergent sequences of automorphisms are preserved.
- If $f_i \to f$ in $\mathrm{Aut}(M)$, then since the cofinite filter is contained in any ultrafilter, $\prod_{i \to \mathcal{U}} f_i$ agrees with $\prod_{i \to \mathcal{U}} f$ over the diagonal copy of M in $M^{\mathcal{U}}$. That is, $(\prod_{i \to \mathcal{U}} f_i) \circ \Delta_M = (\prod_{i \to \mathcal{U}} f) \circ \Delta_M$.
- Applying X and using that X is a Δ-functor, conclude that $\prod_{i \to \mathcal{U}} X(f_i)$ agrees with $\prod_{i \to \mathcal{U}} X(f)$ over the diagonal copy of $X(M)$ inside $X(M)^{\mathcal{U}}$.
- For any point $a \in X(M)$, the above says the sequence $(X(f_i)(a))_{i \in I} =_{\mathcal{U}} (X(f)(a))_{i \in I}$.
- Since $\mathcal{U}$ was arbitrary and the cofinite filter on I is the intersection of all non-principal ultrafilters on I, we conclude that the above equation holds cofinitely. Hence, $X(f_i) \to X(f)$.

Examples, I

Anyways, let's see some more examples of SCC in action.

$\mathrm{Hom}(\mathbb{Q}, -) : \mathbf{Ab} \to \mathbf{Set}$ is not definable (in fact, it's absolutely undefinable! There's no isomorphism to any definable functor.) Why?
Torsion:

$$\prod_{p \to \mathcal{U}} \mathrm{Hom}\left(\mathbb{Q}, \mathbb{Z}/p\mathbb{Z}\right) = 0 \neq \mathrm{Hom}(\mathbb{Q}, \prod_{p \to \mathcal{U}} \mathbb{Z}/p\mathbb{Z})$$

- There's no sensible way to imagine this as living in some eq-sort of the theory, so this lies outside of the scope of Beth's theorem.

Examples, II

"Non-definable" subsets of models can still be definable functors (allowing non-strict transition isomorphisms):

- Let T be the theory of equality expanded by countably many constants.
- If X takes a model M to the even constants in M plus the realizations of the unique non-isolated 1-type, then while viewed as a subfunctor of M it is not an ultrafunctor, it is still isomorphic to M by reindexing the constants.
- That is, after forgetting the data of how it looks like when it's sitting inside the model, X is indistinguishable up to isomorphism by the ultra-stuff from M itself.

Examples, III

- It's obvious that a complete non-isolated type, viewed as a sub-pre-ultrafunctor of its sort, doesn't commute with ultraproducts.
- Let's say that a functor is absolutely undefinable if it not isomorphic to any evaluation functor, i.e. there is no transition isomorphism making it into an ultrafunctor.
- Formally, transition isomorphisms are defined as a collection of bijections $\Phi^{\mathcal{U}}_{(M_i)}$ (ranging over all set-indexed sequences of models and all ultrafilters $\mathcal{U}$ no the indexing sets) such that for any sequence of elementary embeddings $M_i \xrightarrow{f_i} N_i$, the following square commutes:

$$\begin{array}{ccc} X\left(\prod_{i \to \mathcal{U}} M_i\right) & \xrightarrow{\Phi^{\mathcal{U}}_{(M_i)}} & \prod_{i \to \mathcal{U}} X(M_i) \\ {\scriptstyle X\left(\prod_{i\to\mathcal{U}} f_i\right)}\downarrow & & \downarrow{\scriptstyle \prod_{i\to\mathcal{U}} X(f_i)} \\ X\left(\prod_{i \to \mathcal{U}} N_i\right) & \xrightarrow[\Phi^{\mathcal{U}}_{(N_i)}]{} & \prod_{i \to \mathcal{U}} X(N_i). \end{array}$$

Proposition.

Complete non-isolated types are absolutely undefinable.

Proof.

- Transition isomorphisms must be $\text{Aut}(M)^{\mathcal{U}}$-equivariant.
- But, $X(M)^{\mathcal{U}}$ is its own $\text{Aut}(M)^{\mathcal{U}}$-orbit while $X(M^{\mathcal{U}})$ is made up of many $\text{Aut}(M)^{\mathcal{U}}$-orbits.
- So, there can be no transition isomorphism.

□

Examples, IV

We can actually get a lot of mileage out of $\mathrm{Aut}(M)^{\mathcal{U}}$-orbit-counting. For example, let's try to rule out the possibility of infinite disjunctions:

- Counting orbits of $\left(\bigvee_{i \in I} \varphi_i(x_i)\right)(M^{\mathcal{U}})$ gives the sum $\sum_{i \in I} |D_{\varphi_i(x_i)}|^{\mathcal{U}}$, where $D_{\varphi_i(x_i)}$ consists of the complete types containing $\varphi_i(x_i)$.
- Counting orbits of $\left(\bigvee_{i \in I} \varphi_i(x_i)(M)\right)^{\mathcal{U}}$ gives a lower bound $|I^{\mathcal{U}}|$.

Therefore, when situations arise such that

$$\sum_{i \in I} |D_{\varphi_i(x)}|^{\mathcal{U}} < |I^{\mathcal{U}}|,$$

we can conclude that things are absolutely undefinable. For example, what if you're in a theory where all the formulas only contain finitely many types? Hmmmmm...

ω-categorical theories

- A first-order theory T is ω-categorical if it has one countable model up to isomorphism.
- ω-categorical theories have only finitely many types in each sort.
- A theorem of Coquand, Ahlbrandt and Ziegler says that, given two ω-categorical theories T and T' with countable models M and M', a topological isomorphism $\text{Aut}(M) \simeq \text{Aut}(M')$ induces a bi-interpretation $M \simeq M'$.
- Since we know Δ-functors induce continuous maps on automorphism groups, they're a good candidate for definable functors.
- But note that *a priori*, continuity of the automorphism group action is not enough. What about infinite disjoint unions of types?

A definability criterion for ω-categorical theories

Theorem.

(H.) Let $X : \mathbf{Mod}(T) \to \mathbf{Set}$. If T is ω-categorical, the following are equivalent:

1. *For some transition isomorphism, (X, Φ) is a Δ-functor (preserves ultraproducts and diagonal maps).*
2. *For some transition isomorphism, (X, Φ) is definable.*

A definability criterion for ω-categorical theories

Proof.

(Sketch.)

- One direction is immediate by SCC: definable functors are ultrafunctors are at least Δ-functors.
- Let M be the countable model. Use the lemma about Δ-functors (X, Φ) inducing continuous maps on the automorphism groups (equivalently, (X, Φ) has the finite support property) to cover each $\mathrm{Aut}(M)$-orbit of $X(M)$ by a projection from an $\mathrm{Aut}(M)$-orbit of M. By ω-categoricity, the kernel relation of this projection is definable, so we know that $X(M)$ looks like an (*a priori*, possibly infinite) disjoint union of types.
- By orbit-counting (there's also a compactness argument), there are actually only finitely many types.
- Invoke Keisler-Shelah to transfer to all $N \models T$.

□

A definability criterion for ω-categorical theories

Corollary.

Let T and T' be ω-categorical. Let X be an equivalence of categories

$$\mathbf{Mod}(T_1) \overset{X}{\simeq} \mathbf{Mod}(T_2).$$

Then X was induced by a bi-interpretation $T_1 \simeq T_2$ if and only if X was a Δ-functor.

In particular, Bodirsky, Evans, Kompatscher and Pinsker gave an example of two ω-categorical theories T, T' with abstractly isomorphic but not topologically isomorphic automorphism groups of the countable model. This abstract isomorphism induces an equivalence $\mathbf{Mod}(T) \simeq \mathbf{Mod}(T')$ and since it can't come from an interpretation, from the corollary we conclude that it fails to preserve an ultraproduct or a diagonal map was not preserved.

Maybe we can also prove Coquand-Ahlbrandt-Ziegler from this?

Exotic pre-ultrafunctors

In light of the previous result, a natural question to ask is:

Question.

Is being a Δ-functor enough for SCC? That is, do non-definable Δ-functors exist?

Theorem.

(H.) The previous definability criterion fails for general T. That is:

- *There exists a theory T and a Δ-functor $(X, \Phi) : \mathbf{Mod}(T) \to \mathbf{Set}$ which is not definable.*
- *There exists a theory T and a pre-ultrafunctor (X, Φ) which is not a Δ-functor (hence, is also not definable.)*

Exotic pre-ultrafunctors

Proof.

(Sketch.)

- Complete types won't work, so take a complete type and cut it in half into two partial types, one of which refines the other.
- Taking ultraproducts creates external realizations of either one.
- You can either try to construct a transition isomorphism which turns it into a pre-ultrafunctor (creating a non-Δ pre-ultrafunctor) or obtain one non-constructively (creating a non-definable Δ-functor).

□

An internal general adjoint functor theorem

(Since I promised something about an internal general adjoint functor theorem in my abstract, I'll state the result.)

- Does Freyd's general adjoint functor theorem hold for definable adjunctions in any first-order theory?
- The usual proof internalizes except for the part where you need to choose an initial object of the comma categories.
- To do this, you need a choice function for that set. (In model theory, we say that a theory has Skolem functions if **Def**(T) satisfies the external axiom of choice.)

A general adjoint functor theorem internal to **Def**(T)

Definition.

A definable functor $G : \mathbf{D} \to \mathbf{C}$ satisfies the definable solution set condition if there is a uniformly definable family of sets $X_c = \varphi(x, c)$ where each X_c is a weakly initial family in the comma category $\mathrm{Pt}(c, G)$.

Theorem.

(H.) Let $G : D \to C$ be a definable functor between definable categories, such that D has all limits over all definable subcategories. Then:

- *G has a definable left adjoint F with unit of the adjunction also definable, if and only if*
- *G preserves all limits over all deinable subcategories of D, satisfies the definable solution set condition, and the theory T has choice functions for the initial objects of the comma categories $\mathrm{Pt}(c, G)$.*

Thank you!