Introduction to Continuous Model Theory

C. Ward Henson
University of Illinois
www.math.uiuc.edu/~henson/

March 20, 2017
ASL annual meeting, Boise, Idaho
Special Session on Continuous Model Theory

Basic ideas

To transform classical predicate logic to the "continuous" generalization we will emphasize here:

- Replace the space of truth values $\{T, F\}$ by the compact interval $[0, 1]$.
- Replace quantifiers $\forall x$ and $\exists x$ by $\sup_x$ and $\inf_x$.
- Use continuous functions $[0, 1]^n \to [0, 1]$ as connectives.

Non-logical symbols

- A signature $\mathcal{L}$ consists of function and predicate symbols, as usual.
  - n-ary function symbols: interpreted as functions $M^n \to M$.
  - n-ary predicate symbols: interpreted as functions $M^n \to [0, 1]$.
- The metric is considered as a (logical) predicate (analogous to the use of equality in classical logic).
- $\mathcal{L}$ specifies a modulus of uniform continuity for each function symbol and predicate symbol. (e.g.: 1-Lipshitz.)
- $\mathcal{L}$-terms and atomic $\mathcal{L}$-formulas are built inductively as in classical logic.

Connectives and Quantifiers

- Any continuous function $[0,1]^n \to [0,1]$ is admitted as an n-ary connective.
  - That makes syntax uncountable. BUT, a uniformly dense set of connectives is good enough. Hence it suffices to work with a countable set of connectives, by the lattice version of the Stone-Weierstrass Theorem.
- Use "$\sup_x \varphi$" and "$\inf_x \varphi$" as quantifiers in place of "$\forall x \varphi$" and "$\exists x \varphi$".

Notation: $s \dot{-} t := \max(s - t, 0)$ and $s \dot{+} t := \min(s + t, 1)$.

Structures

Definition

An $\mathcal{L}$-pre-structure is a set M, equipped with a pseudo-metric d^M and uniformly continuous (as specified by $\mathcal{L}$) interpretations f^M, P^M of all symbols $f, P \in \mathcal{L}$.

M is an $\mathcal{L}$-structure if d^M is a complete metric.

Semantics

As usual, the notation $\varphi(x_1, \ldots, x_n)$ [or simply $\varphi(x)$] means that the free variables of φ are among $x_1, \ldots, x_n$, which are distinct.

If M is a pre-structure and $a \in M^n$, we define the value $\varphi^M(a) \in [0, 1]$ inductively, in the "obvious way".

- $\big(P(t_1(x), \ldots, t_k(x))\big)^{\mathcal{M}}(a) = P^{\mathcal{M}}(t_1^{\mathcal{M}}(a), \ldots, t_k^{\mathcal{M}}(a))$;
- $\big(u(\varphi_1(x), \ldots, \varphi_k(x))\big)^{\mathcal{M}}(a) = u(\varphi_1^{\mathcal{M}}(a), \ldots, \varphi_k^{\mathcal{M}}(a))$;
- $\big(\sup_y \varphi(x, y)\big)^{\mathcal{M}}(a) = \sup\{\varphi^{\mathcal{M}}(a, b) \mid b \in M\}$.
- $\big(\inf_y \varphi(x, y)\big)^{\mathcal{M}}(a) = \inf\{\varphi^{\mathcal{M}}(a, b) \mid b \in M\}$.

Lemma

Each function $\varphi^M \colon M^n \to [0,1]$ is uniformly continuous, and the moduli of uniform continuity only depend on the signature of M.

Proof.

By induction on φ. $\square$

Completion of a pre-structure

Suppose M is a prestructure.

- An equivalence relation $\sim_d$ may be defined on M by $a \sim_d b :\Leftrightarrow d(a,b) = 0$. If $a \sim_d a'$ and $b \sim_d b'$, we have $d(a,b) = d(a',b')$. So we may define a metric on the quotient $M/\sim_d$ by $d([a],[b]) = d(a,b)$.
- The interpretations of predicate and function symbols in M are uniformly continuous. Hence they have canonical extensions to the completion of the quotient metric space $M/\sim_d$.
- We denote the resulting structure by $\widehat{M/\sim_d}$ and call it the completion of the pre-structure M.

Each prestructure M is logically indistinguishable from its completion:

Fact

The canonical quotient map $\pi\colon M \to \widehat{M/{\sim_d}}$ is elementary (i.e., π preserves the value of all formulas).

Ultraproducts

- $(M_i \mid i \in I)$ are $\mathcal{L}$-pre-structures, $\mathcal{U}$ an ultrafilter on I.
- We let $N_0 = \prod_{i \in I} M_i$ as a set; its members are $(a) = (a_i \mid i \in I)$, $a_i \in M_i$.
- We interpret the symbols:

$$\begin{aligned} f^{N_0}\big((a_i \mid i \in I), \dots\big) &= \big(f^{M_i}(a_i, \dots) \mid i \in I\big) && \in N_0 \\ P^{N_0}\big((a_i \mid i \in I), \dots\big) &= \lim_{i,\mathcal{U}} P^{M_i}(a_i, \dots) && \in [0,1] \end{aligned}$$

- This way N_0 is a pre-structure. We define the ultraproduct to be $\widehat{N_0}$ (the completion) and denote it by $\prod_{i \in I} M_i / \mathcal{U}$.
- The image of $(a) \in N_0$ in $\widehat{N_0}$ is denoted $(a)_{\mathcal{U}}$; note that

$$(a)_{\mathcal{U}} = (b)_{\mathcal{U}} \iff 0 = \lim_{i,\mathcal{U}} d(a_i, b_i) \quad \Big[= d^{N_0}\big((a),(b)\big)\Big].$$

Properties of ultraproducts

- Łoś's Theorem: for every formula $\varphi(x, y, \ldots)$ and elements $(a)_{\mathcal{U}}, (b)_{\mathcal{U}}, \ldots \in N = \prod M_i/\mathcal{U}$:

$$\varphi^N((a)_{\mathcal{U}}, (b)_{\mathcal{U}}, \ldots) = \lim_{i,\mathcal{U}} \varphi^{M_i}(a_i, b_i, \ldots).$$

- [Easy] $M \equiv N$ (M and N are elementarily equivalent) if and only if M admits an elementary embedding into an ultrapower of N.
- [Deeper: generalising Keisler & Shelah] $M \equiv N$ if and only if M and N have ultrapowers that are isomorphic.

Types (without parameters)

Definition

Let M be an $\mathcal{L}$-structure, $a \in M^n$. Then:

$$\mathrm{tp}^M(a) = \{\varphi(x) \mid \varphi(x) \in \mathcal{L}, \varphi(a)^M = 0\} \equiv \Big[\mathrm{Th}(M, a)\Big].$$

$\mathrm{S}_n(T)$ is the space of types of n-tuples in models of T.

If $p \in \mathrm{S}_n(T)$ and $\varphi(x)$ is a formula, we may define the value of $\varphi(x)$ at (a realization of) p, denoted $\varphi(x)^p$ as follows:

$$\varphi(x)^p = r \quad :\Longleftrightarrow \quad \big|\varphi(x) - r\big| \in p.$$

- The (logic) topology on $S_n(T)$ is the minimal topology such that $p \mapsto \varphi^p$ is continuous for all φ.

Types (with parameters)

Definition

Let M be a structure, $a \in M^n$, $B \subseteq M$. Then:

$$\operatorname{tp}^M(a/B) = \{\varphi(x,b) \mid \varphi(x,y) \in \mathcal{L}, b \in B^m, \varphi(a,b)^M = 0\}.$$

$\mathrm{S}_n(B)$ is the space of types over B of n-tuples in elementary extensions of M. If $p \in \mathrm{S}_n(B)$, $b \in B$:

$$\varphi(x,b)^p = r \quad :\Longleftrightarrow \quad \big|\varphi(x,b) - r\big| \in p.$$

The logic topology on $S_n(B)$ is minimal such that $p \mapsto \varphi(x,b)^p$ is continuous for all $\varphi(x,b)$, $b \in B^m$. It is compact and Hausdorff.

Definable predicates

- We identify a formula $\varphi(x)$ with the (continuous) function $\widetilde{\varphi}\colon \mathrm{S}_n(T) \to [0,1]$ it induces: $p \mapsto \varphi^p$. By Stone-Weierstrass these functions are dense in $C(\mathrm{S}_n(T),[0,1])$.
- An arbitrary continuous function $P\colon \mathrm{S}_n(T) \to [0,1]$ is called a definable predicate. It is a uniform limit of functions induced by formulas: $P = \lim_{k\to\infty} \widetilde{\varphi_k}$. Its interpretation in $M \vDash T$ is given for $a \in M^n$ by:

$$P^M(a) := P(\mathrm{tp}_M(a)) = \lim_k \varphi_k^M(a).$$

  Since each φ_k^M is uniformly continuous, so is P^M.
- The same applies with parameters. Note that a definable predicate $\lim \varphi_k(x, b_k)$ over an infinite set B of parameters may depend on a countably infinite sequence of elements of B.

Metric on types

The topological structure of $\mathrm{S}_n(T)$ is insufficient. We also need to consider the distance between types:

$$d(p, q) = \inf\{d(a, b) \mid a, b \in M \vDash T \text{ and } M \vDash p(a) \cup q(b)\}.$$

(In case T is incomplete and p, q belong to different completions: $d(p, q) = \inf \varnothing := \infty$.)

Some properties of $S_n(T)$

Type spaces in continuous model theory are equipped with a topology (*i.e.* the “logic topology”) and a metric (induced by the metric on models of T).

The topology is compact, Hausdorff; the metric is complete, and its topology is finer.

Convention: unmodified topological concepts (continuous, closed, open, dense, . . .) always refer to the (logic) topology; metric concepts (uniform continuity, uniform convergence, . . .) always refer to the (induced) metric. Exceptions to this convention are made explicit.

Facts

- *Every continuous function* $\mathrm{S}_n(T) \to [0,1]$ *is uniformly continuous.*
- *The metric* $d\colon \mathrm{S}_n(T)^2 \to [0,1]$ *is lower semi-continuous. That is, for all* $r \in [0,1]$, *the set*
$$\{(p,q) \in \mathrm{S}_n(T)^2 \mid d(p,q) \leq r\}$$
is closed in $\mathrm{S}(T)^2$.

Definable sets

Basic setting: T is a theory and we have a definable predicate P (of n arguments) in models of T. Recall this means we have a continuous function $P\colon \mathrm{S}_n(T) \to [0,1]$ whose interpretation in any $M \vDash T$ is given (for each $a \in M^n$) by

$$P^M(a) = P(\mathrm{tp}_M(a))$$

Equivalently, we have a sequence $(\varphi_n(x)) = (\varphi_{P,n}(x))$ of formulas such that (for each $M \vDash T$ and each $a \in M^n$)

$$P^M(a) = \lim_n \varphi_n^M(a)$$

and we require that the convergence in this limit is uniform in the tuple a and in the model M.

For such a situation, we let Z_P^M to be the zeroset of P^M, whenever $M \models T$:

$$Z_P^M = \{a \in M^n \mid P^M(a) = 0\}$$

Note that Z_P^M is always a type-definable set; indeed for a suitable sequence (ϵ_n) of reals, we have

$$Z_P^M = \bigcap_n \{a \in M^n \mid \varphi_n^M(a) \leq \epsilon_n\}$$

for every $M \models T$.

Moreover, for any $M \models T$ and $a \in M^n$ we see that $a \in Z_P^M$ if and only if $\mathrm{tp}_M(a)$ is in the closed subset of $\mathrm{S}_n(T)$ given by

$$\pi_P^{-1}(0) = \bigcap_n \{q \in \mathrm{S}_n(T) \mid \varphi_n(x) \leq \epsilon_n \text{ is an element of } q\}$$

which we will denote by $[P = 0]$.

Definition

We say the zerosets of P are definable in models of T if the predicate on M^n giving the distance to Z_P^M is a definable predicate in models M of T.

That is, there should be a definable predicate Q in models of T such that for every $M \vDash T$ and every $a \in M^n$ we have

$$Q^M(a) = \mathrm{dist}(a, Z_P^M)$$

Lemma

The zerosets of P are definable in models of T if and only if the function (from $S_n(T)$ to $[0,1]$) given by

$$dist(q, [P = 0])$$

is continuous.

Proof: For every $M \models T$ and $a \in M^n$ we have

$$\mathrm{dist}(a, Z_P^M) = \mathrm{dist}(tp_M(a), [P = 0])$$

The following result explains why this is the "right" definition of "definable set" in the model theory of metric structures:

Theorem

Let P be an n-ary definable predicate in models of T. The following are equivalent.
(1) The zerosets of P are definable in models of T,
(2) For any $(m+n)$-ary definable predicate Q in models of T, there is an m-ary definable predicate R in models of T such that for any $M \models T$ and any $a \in M^m$ we have

$$R^M(a) = \inf\{Q^M(a,b) \mid b \in M^n \text{ and } P^M(b) = 0\}$$

That is, we can "quantify" over the zeroset of P without leaving the category of definable predicates.

Proof: ($\Leftarrow$) Apply condition 2 to the formula $d(x, y)$.
($\Rightarrow$) By uniform continuity of Q^M and a real analysis argument there exists a continuous function $\alpha \colon [0, 1] \to [0, 1]$ with $\alpha(0) = 0$ such that

$$|Q^M(a, b) - Q^M(a, c)| \leq \alpha(d(b, c))$$

for all $a \in M^m$ and all $b, c \in M^n$; moreover, α can be taken independent of M. Now check that

$$\inf\{Q^M(a, b) \mid b \in Z_P^M\} = \inf\{Q^M(a, c) + \alpha(\mathrm{dist}(c, Z_P^M)) \mid c \in M^n\}$$

and note that the right side is a definable predicate when condition (1) holds.

Criterion 1 for definability of zerosets

Proposition

The following are equivalent for a definable predicate P in models of T.
(1) The zerosets of P are definable sets in models of T.
(2) For all $\epsilon > 0$ there exists $\delta > 0$ such that for all $M \models T$ and all $a \in M^n$ we have

$$P^M(a) < \delta \text{ implies } \exists b \in M^n (d(a,b) \leq \epsilon \text{ and } P^M(b) = 0)$$

Condition (2) says that the sets $\{a \in M^n \mid P^M(a) < \delta\}$ Hausdorff converge to the zeroset of P^M in M^n. Note that condition (1) depends only on the zerosets of P and not on P otherwise.

Proof ($\Rightarrow$) Let Q be the definable predicate with interpretations $\text{dist}(a, Z_P^M)$ for all $M \models T$. Note that in every model, P^M and Q^M have the same zeroset. Using a compactness argument, we see that for every $\epsilon > 0$ there exists $\delta > 0$ so that for all $M \models T$ and all $a \in M^n$ we have

$$P^M(a) \leq \delta \text{ implies } Q^M(a) < \epsilon$$

This implies condition (2).
($\Leftarrow$) Assuming condition (2) we get a continous function $\alpha\colon [0,1] \to [0,1]$ with $\alpha(0) = 0$ and $\text{dist}(a, Z_P^M) \leq \alpha(P^M(a))$ for all $M \models T$ and all $a \in M^n$. Then let Q(x) be the definable predicate given by $\inf_y(\alpha(P(y)) \dot{+} d(x,y))$ and show that its interpretation is always equal to $\text{dist}(a, Z_P^M)$.

Criterion 2 for definability of zerosets

Proposition

The following are equivalent for a definable predicate P in models of T.
(1) The zerosets of P are definable sets in models of T.
(2) For any family $(M_i \mid i \in I)$ of models of T and any ultrafilter $\mathcal{U}$ on I, with N the $\mathcal{U}$-ultraproduct of the family, every member of Z_P^N is represented by a family $(a_i \mid i \in I)$ with $a_i \in Z_P^{M_i}$ for a $\mathcal{U}$-large set of I.

Cautionary example 1

Let T be the theory of pointed metric spaces of diameter ≤ 1 in which the metric satisfies the ultrametric inequality:

$$d(x, y) \leq \max(d(x, z), d(z, y))$$

(in a signature with c the constant symbol naming the distinguished element) and let $r \in (0, 1)$. The closed balls of radius r centered at c are not definable in models of T.

Note that the closed ball of radius r is the zeroset of the formula $P(x) = (d(c, x) \dot{-} r)$. It is not hard to show that P fails to satisfy Criterion 1.

Cautionary example 2

Theorem

The following are equivalent for a complete theory T in a countable signature.
(1) T is separably categorical.
(2) Every zeroset of a definable predicate in T is a definable set.

Proof: $(\Rightarrow)$ Let X be a closed set in $\mathrm{S}_n(T)$. The function $\mathrm{dist}(q, X)$ is obviously continuous for the metric topology, so it is continuous since T is separably categorical.
$(\Leftarrow)$ Take any $p \in \mathrm{S}_n(T)$ and consider the closed set $X = \{p\}$. Condition (2) implies that $d(p, q)$ is a continuous function of q on $\mathrm{S}_n(T)$, from which it follows that p is isolated. Since n and p were arbitrary, this implies T is separably categorical.

Now suppose T is a theory in a countable signature and that M is a separable but noncompact model of T. Let $(a_n \mid n \geq 1)$ be an enumeration of a dense set in M, and let

$$T' = \mathrm{Th}(M, (a_n)_{n\geq 1}).$$

Then T' is not separably categorical. Hence by the previous argument there is a predicate P definable in T' whose zerosets are not definable over T'.

This shows that zerosets that are not definable sets occur in every theory of interest (when parameters are allowed).

Imaginaries

Let T be a theory in signature $\mathcal{L}$. It is convenient to introduce a many-sorted signature $\mathcal{L}^{eq}$ expanding $\mathcal{L}$ together with an $\mathcal{L}^{eq}$-theory T^{eq} containing T.

Our purpose is to add sorts of imaginaries over which we can then quantify. We do this without increasing the expressive power of our language, as each model of T expands to an essentially unique model of T^{eq}.

Definability in models of T^{eq} corresponds exactly to interpretability in models of T.

$(\mathcal{L}^{eq}, T^{eq})$ is then taken to be the smallest expansion of $(\mathcal{L}, T)$ that is closed under the following three operations:

Operation 1: countable products

Setting: $(S_i \mid i \in I)$ is a finite or countable family of sorts, with $I = \{1, 2, 3, \ldots\}$.

We introduce a new sort U, a metric symbol d_U, and a unary function $\pi_i \colon U \to S_i$ for each $i \in I$.

Given any model M of T, we interpret U as the product $\prod_{i \in I} S_i^M$ and take $\pi_i^M \colon U^M \to S_i^M$ to be the coordinate projection. We interpret d_U to be the metric defined for $a = (a_i), b = (b_i)$ by

$$d_U^M(a, b) := \sum_{i \in I} 2^{-i} d_{S_i}(a_i, b_i).$$

The modulus of uniform continuity for π_i is taken to be $\Delta_i(\epsilon) = 2^{-i}\epsilon$.

Further, we add axioms to T expressing the following statements:

- for each $i \in I$ and every $a, b \in U$

$$d_U(a,b) \le 2^{-i} + \sum_{j=1}^{i} 2^{-i} d_{S_i}(\pi_i(a), \pi_i(b)).$$

- for each $i \in I$ and all $a_j \in S_j, j = 1, \ldots, i$

$$\inf_{x \in U} \max_{1 \le j \le i} \left(|\pi_j(x) - a_j| \right) = 0$$

These ensure that the map (π_i) from U into $\prod_{i \in I} S_i$ is distance preserving and surjective.

Operation 2: definable sets

Setting: S is a sort and P is a definable predicate on S whose zeroset is a definable set.

We introduce a new sort U, a metric symbol d_U, and a unary function symbol $\eta\colon U \to S$.

Given any model M of T, we interpret U as the zeroset of P in M, interpret d_U as the restriction of d_S, and interpret η as the inclusion.

The modulus of uniform continuity for η is taken to be $\Delta(\epsilon) = \epsilon$.

Further, we add axioms to T expressing the following statements:

- for every $a, b \in U$

$$d_S(\eta(a), \eta(b)) = d_U(a, b) \quad \text{and} \quad P(\eta(a)) = 0.$$

- for every c in the zeroset of P

$$\inf_{x \in U} d_s(c, \eta(x)) = 0.$$

These ensure that the map η is distance preserving and maps onto the zeroset of P.

Operation 3: quotients

Setting: S is a sort and $D(x, y)$ is a binary definable predicate on S such that for every model M of T, the function D^M is a pseudo-metric on S^M.

We introduce a new sort U, a metric symbol d_U, and a unary function symbol $\pi \colon S \to U$.

Given any model of T, we interpret (U, d_U) as the completion of the quotient metric space of (S^M, D^M) and interpret π as the quotient map from S^M into U.

The modulus of uniform continuity for π is taken to be the same as the modulus of the definable predicate D.

Further, we add axioms to T expressing the following statements:

- for every $a, b \in S$

$$d_U(\pi(a), \pi(b)) = D(a, b).$$

- for every $c \in U$

$$\inf_{x \in S} d_U(c, \pi(x)) = 0.$$

These ensure that (U, d_U, π) is isometrically isomorphic to the completion of the quotient metric space of (S, D). (Given $a \in S$, map $[a]_D$ to $\pi(a) \in U$ and extend to the completion.)

Some further topics

On the following slides (not used in the ASL special session talk) we mention some other aspects of continuous model theory that are of importance, especially in application areas.

- Origins of continuous model theory
- Elementary notions.
- Quantifier elimination
- Existentially closed models, etc
- Omitting types and separable categoricity

Finally, on the last pages, we give a few references.

Origins

Many classes of metric structures arising in analysis and geometry have been seen to be well-behaved model theoretically, although they are not elementary in the classical sense.
Continuous logic is an attempt to apply model-theoretic tools smoothly and effectively to such classes. It's main antecedents are:

- Use of an ultraproduct construction in analysis and similar constructions in nonstandard analysis (starting in the 1960s).
- Henson's logic for Banach structures (positive bounded formulas, approximate satisfaction) (starting in the 1970s).
- Ben-Yaacov's positive logic and compact abstract theories (starting about 2000).
- Krivine's real-valued logic (only used universal quantifiers) (1974).

Earlier related work includes:

- Łukasiewicz's $[0,1]$-valued propositional logic (no quantifiers, only used special connectives; starting about 1930).
- Chang and Keisler's book on continuous model theory (1966).
- ...perhaps others ...

Various “elementary” notions

- Elementary equivalence (denoted $M \equiv N$): If M, N are two pre-structures then $M \equiv N$ if $\varphi^M = \varphi^N \in [0,1]$ for every sentence φ (i.e.: formula without free variables).
  Equivalently: “$\varphi^M = 0 \iff \varphi^N = 0$ for all sentences φ.”
- Elementary extension (denoted $M \preceq N$): This holds if $M \subseteq N$ and $\varphi^M(a) = \varphi^N(a)$ for every formula φ and $a \in M$. It implies $M \equiv N$.

Lemma (Elementary chains)

The union of an elementary chain $M_0 \preceq M_1 \preceq \ldots$ is an elementary extension of each M_i.

Caution: by the union of an increasing chain we mean the completion of its set-theoretic union.

- A theory T is a set of sentences (closed formulas).
- M is a model of T (written $M \vDash T$) $\iff \varphi^M = 0$ for all $\varphi \in T$ AND M is a structure.
- We sometimes write T as a set of conditions "$\varphi = 0$". We may also allow as conditions things of the form "$\varphi \leq r$", "$\varphi \geq r$", "$\varphi = r$", etc.
- If M is any structure then its theory is

$$\mathrm{Th}(M) = \{\varphi \mid \varphi^M = 0\} \quad \Big[\equiv \{\text{“}\psi = r\text{”} \mid \psi^M = r\} \Big].$$

Theories of this form are called complete (equivalently: complete theories are the maximal satisfiable theories).

Theorem (Compactness)

A theory is satisfiable if and only if it is finitely satisfiable.

Note that:

$$T \equiv \{\text{“}\varphi \leq 2^{-n}\text{”} \mid n < \omega \text{ and } \text{“}\varphi = 0\text{”} \in T\}.$$

Corollary

Assume that T is approximately finitely satisfiable. Then T is satisfiable.

Notation: $\|\cdot\|$ denotes the metric density character.
$\mathcal{L}$ is a signature with $\leq \kappa$ symbols.
M is an $\mathcal{L}$-structure and $A \subseteq M$ has $\|A\| \leq \kappa$.

Theorem (Löwenheim-Skolem Theorem)

There exists an elementary substructure N of M such that $A \subseteq N$ and $\|N\| \leq \kappa$.

Elementary (axiomatizable) classes

A class $\mathcal{C}$ of $\mathcal{L}$-structures is elementary or axiomatizable if there is an $\mathcal{L}$-theory T such that $\mathcal{C}$ is the class of all models of T. When this holds we call T a set of axioms for $\mathcal{C}$.

Theorem

Let $\mathcal{C}$ be a class of $\mathcal{L}$-structures. Then $\mathcal{C}$ is axiomatizable iff $\mathcal{C}$ is closed under isomorphisms, ultraproducts, and ultraroots.

Here: M is an ultraroot of N if N is isomorphic to some ultrapower of M.

Definition

Let κ be a cardinal, M a structure.

- M is κ-saturated if for every $A \subseteq M$ such that $|A| < \kappa$ and every $p \in S_1(A)$: p is realized in M.
- M is κ-homogeneous if for every $A \subseteq M$ such that $|A| < \kappa$ and every mapping $f : A \to M$ that preserves truth values, f extends to an automorphism of M.

Fact

Let M be any structure and $\mathcal{U}$ a non-principal ultrafilter on $\mathbb{N}$. Then the ultrapower $M^{\aleph_0}/\mathcal{U}$ is ω_1-saturated.

A monster model of a complete theory T is a model of T that is κ-saturated and κ-homogeneous for some κ that is much larger than any set under consideration.

Fact

- *Every complete theory T has a monster model.*
- *If $\mathbb{M}$ is a monster model for T, then every "small" model of T (i.e., smaller than κ) is isomorphic to some $N \preceq \mathbb{M}$.*
- *If $A \subseteq \mathbb{M}$ is small then $\mathrm{S}_n(A)$ is the set of orbits in $\mathbb{M}^n$ under $\mathrm{Aut}(\mathbb{M}/A)$.*

Thus monster models serve as "universal domains": everything happens inside, and the automorphism group is large enough.

Quantifier elimination

Suppose $M, N \vDash T$ and $a \in M^n, b \in N^n$. We write

$$a \equiv_0 b \text{ in } M, N$$

if $\varphi^M(a) = \varphi^N(b)$ for every quantifier-free formula $\varphi(x)$.

Note that this is equivalent to the existence of an isomorphism J from $\langle a \rangle_M$ onto $\langle b \rangle_N$ satisfying $J(a_i) = b_i$ for all $i = 1, \ldots, n$.

Likewise we write

$$a \equiv b \text{ in } M, N$$

if $\varphi^M(a) = \varphi^N(b)$ for every formula $\varphi(x)$, which is the same as saying

$$(M, a) \equiv (N, b)$$

T admits quantifier elimination if

- for every formula $\varphi(x)$ and every $\epsilon > 0$ there is a quantifier-free formula such that the following condition holds in all models of T:

$$\sup_x \left|\varphi(x) - \psi(x)\right| \leq \epsilon$$

Here $x = x_1, \ldots, x_n$ is a finite tuple and $\sup_x$ means $\sup_{x_1} \cdots \sup_{x_n}$.

Criterion 1 for QE: types = qf types

T admits QE if and only if

- for all $n \geq 1$ and all $p \in S_n(T)$, the type p is determined by the quantifier free formulas it contains.

Proof: ($\Leftarrow$) Consider the restriction map $\pi\colon S_n(T) \to S_n^{qf}(T)$. This map is always continuous and surjective. QE is equivalent to its being a topological homeomorphism. Since both spaces are compact, for QE it suffices that the map is injective.

Criterion 2 for QE: back-and-forth property

T admits QE if and only if

- for every ω-saturated models M, N of T and every $a \in M^n, b \in M, a' \in N^n$, if $a \equiv_0 a'$ in M, N, then there exists $b' \in N$ such that $(a, b) \equiv_0 (a', b')$ in M, N.

Proof: $(\Leftarrow)$ Use the condition to show that whenever $a \equiv_0 a'$ in M, N, then a and a' are given the same values by formulas of the form $\inf_y \varphi(x, y)$ in which φ is quantifier-free. Use an argument like that for Criterion 1 to show that every formula of this kind is approximated uniformly by quantifier-free formulas. Then use induction on syntactic complexity to show the same is true for all formulas.

Criterion 3 for QE: extension of embeddings

T admits QE if and only if

- for every $M, N \models T$, every embedding of a substructure of M into N can be extended to an embedding of M into an elementary extension of N.

Moreover, if the signature of T has $\leq \kappa$ many symbols, then it suffices to consider M of density character $\leq \kappa$ and to consider a fixed elementary extension of N that is κ-saturated.

Proof: ($\Leftarrow$) Show T verifies Criterion 2. So consider ω-saturated models M, N of T and tuples $a \in M^n$, $a' \in N^n$ and an element $b \in M$. Assume $a \equiv_0 a'$ in M, N. This gives a canonical isomorphism J of $\langle a \rangle_M \subseteq M$ onto $\langle b \rangle_N \subseteq N$ that maps a exactly to b. The condition above gives an extension of J to an embedding (which we still call J) of M into an elementary extension N' of N. This ensures $(a, b) \equiv_0 (a', J(b))$ in M, N'. Since N is ω-saturated, we can find $b' \in N$ such that $(a', J(b)) \equiv_0 (a', b')$ in N, N'. Therefore $(a, b) \equiv_0 (a', b')$ in M, N as desired.

Existentially closed models; model companions

$M \models T$ is an existentially closed (e.c.) model of T if

- for every $M \subseteq N \models T$, every quantifier-free formula $\varphi(x, y)$, and every $a \in M^m$, $\quad \big(\inf_y \varphi(a, y)\big)^M = \big(\inf_y \varphi(a, y)\big)^N$.

Here $y = y_1, \ldots, y_n$ is a finite tuple and $\inf_y$ means $\inf_{y_1} \cdots \inf_{y_n}$.

The condition above is equivalent to requiring the implication

$$\big(\inf_y \varphi(a, y)\big)^N = 0 \text{ implies } \big(\inf_y \varphi(a, y)\big)^M = 0$$

for all quantifier-free φ.

- T is model complete if every embedding between models of T is an elementary embedding.
- (Note: quantifier elimination $\implies$ model complete.)

- T^* is a model companion of T if they have the same signature, T^* is model complete, and T, T^* have the same substructures of models.

- T is inductive if its class of models is closed under unions of arbitrary chains.

Theorem

Suppose T is an inductive theory and $M \models T$. Then there exists an e.c. model N of T such that $M \subseteq N$.

Theorem

Let T be an inductive theory and let T^* be any theory with the same signature as T.
(a) T^* is a model companion of T if and only if the models of T^* are exactly the e.c. models of T.
(b) In particular, T has a model companion if and only if the class of e.c. models of T is axiomatizable.

Theorem

Suppose T^* is a model companion of T. If models of T have the amalgamation property over substructures, then T^* admits quantifier elimination.

Note: in the setting of this Theorem, T^* is called the model completion of T.

Separable models; special properties of models

Throughout this subsection on separable models, we take T to be a complete theory with a countable signature.

isolated types, atomic models

- A type $p \in S_n(T)$ is isolated (i.e., principal) if it has the same filter of neighborhoods in the logic topology as in the metric topology.
- A structure M is atomic if every type realized in M is isolated (w.r.t. $T = \mathrm{Th}(M)$).

Theorem (part of the Omitting Types Theorem)

Let $p \in S_n(T)$. The following are equivalent:
(1) p is isolated.
(2) p is realized in every model of T.

Theorem (Atomic models)

(a) T has an atomic model iff for each $n \in \mathbb{N}$, the isolated types in $\mathrm{S}_n(T)$ are dense.

(b) T has at most one separable atomic model.

(c) If $\mathrm{S}_n(T)$ is separable in the metric topology, for all $n \in \mathbb{N}$, then T has a separable atomic model.

- M is ω-near-homogeneous if for any $a, b \in M^n$ realizing the same n-type in M, and any $\epsilon > 0$, there is an automorphism σ of M such that $d(\sigma(a), b) < \epsilon$.

Theorem

If M is a separable atomic model of T, then M is ω-near-homogeneous. In particular, this is true of the unique separable model of a separably categorical theory.

Caution: in discrete structures (classical model theory), the countable atomic model is always ω-homogeneous. In the setting of metric structures, however, ω-near-homogeneity is often the most one has.

Theorem (Separable categoricity)

The following are equivalent:
(1) T has exactly one separable model.
(2) Every separable model of T is atomic.
(3) Every type in $S_n(T)$ is isolated, for all $n \in \mathbb{N}$.
(4) The metric topology is compact on $\mathrm{S}_n(T)$, for all $n \in \mathbb{N}$.
(5) The logic topology agrees with the metric topology on $\mathrm{S}_n(T)$, for all $n \in \mathbb{N}$.

Corollary

Suppose T is separably categorical and T' is the restriction of T to a smaller signature. Then T' is also separably categorical.

Note: this is true in particular if T' is the theory of all metric spaces that arise from models of T. *A priori* T' might have separable models that do not come from any model of T.

Proof: for each n, consider the restriction map from $S_n(T)$ onto $S_n(T')$; it is continuous, contractive, and surjective. Hence it preserves compactness with respect to the metric topologies.

Some references

Itaï Ben Yaacov, Alexander Berenstein, C. Ward Henson, and Alexander Usvyatsov, *Model theory for metric structures*, in Lecture Notes series of the London Mathematical Society, No. 350, Cambridge University Press, 2008, 315–427.
[Gives a general treatment of continuous model theory.]

Itaï Ben Yaacov and Alexander Usvyatsov, *Continuous first order logic and local stability*, Transactions of the American Mathematical Society 362 (2010), 5213–5259.
[Gives a general treatment, with local stability.]

Farah, Hart, Lupini, Robert, Tikuisis, Vignati, Winter, Model theory of C*-algebras, preprint (2016). 141 pages. arXiv:1602.08072v3.
[While focused on C*-algebras, this monograph presents many basic elements of continuous model theory and gives illuminating examples.]

Itaï Ben Yaacov and Alexander Usvyatsov, On d-finiteness in continuous structures, Fundamenta Mathematicae 194 (2007), 67–88.
[Section 1 gives an exposition of theorems about separable models; *e.g.*, the full omitting types theorem, existence of atomic and approx. saturated separable models.]

Earlier papers: compact abstract theories

- Itaï Ben Yaacov, Positive model theory and compact abstract theories, Journal of Mathematical Logic 3, 2003, 85–118.
- Itaï Ben Yaacov, Simplicity in compact abstract theories, Journal of Mathematical Logic 3, 2003, 163–191.
- Itaï Ben Yaacov, Thickness, and a categoric view of type-space functors, Fundamenta Mathematicae, 179, 2003, 199–224.
- Itaï Ben Yaacov, Compactness and independence in non first order frameworks, Bulletin of Symbolic Logic, 11, 2005, 28–50.

Earlier papers: positive bounded formulas

- C. Ward Henson, Nonstandard hulls of Banach spaces, Israel Journal of Mathematics 25, 1976, 108–144.
- C. Ward Henson and L. C. Moore, Nonstandard analysis and the theory of Banach spaces, in Nonstandard Analysis – Recent Developments (Victoria, B.C., 1980), Springer Lecture Notes in Mathematics No. 983, 1983, 27–112.
- C. Ward Henson and José Iovino, Ultraproducts in analysis, in Analysis and Logic, London Mathematical Society Lecture Notes Series, vol. 262, 2002, 1–113.

Earlier papers: Ultraproducts in functional analysis

- Jean Bretagnolle, Didier Dacunha-Castelle, Jean-Louis Krivine, Lois stables et espaces L^p, Annals Inst. Henri Poincaré Sect. B (N.S.) 2, 1965/1966, 231–259.
- D. Dacunha-Castelle and Jean-Louis Krivine, Applications des ultraproduits à l'étude des espaces et des algèbres de Banach, Studia Math. 41, 1972, 315–334.
- Stefan Heinrich, Ultraproducts in Banach space theory, Journal für die Reine und Angewandte Mathathematik 313, 1980, 72–104.