Model theory for metric structures

Itaï Ben Yaacov

Alexander Berenstein

C. Ward Henson

Alexander Usvyatsov

Contents

1 Introduction

A metric structure is a many-sorted structure in which each sort is a complete metric space of finite diameter. Additionally, the structure consists of some distinguished elements as well as some functions (of several variables) (a) between sorts and (b) from sorts to bounded subsets of $\mathbb{R}$, and these functions are all required to be uniformly continuous. Examples arise throughout mathematics, especially in analysis and geometry. They include metric spaces themselves, measure algebras, asymptotic cones of finitely generated groups, and structures based on Banach spaces (where one takes the sorts to be balls), including Banach lattices, C*-algebras, etc.

The usual first-order logic does not work very well for such structures, and several good alternatives have been developed. One alternative is the logic of *positive bounded formulas with an approximate semantics* (see [23, 25, 24]). This was developed for structures from functional analysis that are based on Banach spaces; it is easily adapted to the more general metric structure setting that is considered here. Another successful alternative is the setting of *compact abstract theories* (cats; see [1, 3, 4]). A recent development is the realization that for metric structures the frameworks of positive bounded formulas and of cats are equivalent. (The full cat framework is more general.) Further, out of this discovery has come a new *continuous* version of first-order logic that is suitable for metric structures; it is equivalent to both the positive bounded and cat approaches, but has many advantages over them.

The logic for metric structures that we describe here fits into the framework of continuous logics that was studied extensively in the 1960s and then dropped (see [12]). In that work, any compact Hausdorff space X was allowed as the set of truth values for a logic. This turned out to be too general for a completely successful theory.

We take the space X of truth values to be a closed, bounded interval of real numbers, with the order topology. It is sufficient to focus on the case where X is $[0, 1]$. In [12], a wide variety of quantifiers was allowed and studied. Since our truth value set carries a natural complete linear ordering, there are two canonical quantifiers that clearly deserve special attention; these are the operations sup and inf, and it happens that these are the only quantifiers we need to consider in the setting of continuous logic and metric structures.

The continuous logic developed here is strikingly parallel to the usual first-order logic, once one enlarges the set of possible truth values from $\{0,1\}$ to $[0,1]$. Predicates, including the equality relation, become functions from the underlying set A of a mathematical structure into the interval $[0,1]$. Indeed, the natural $[0,1]$-valued counterpart of the equality predicate is a metric d on A (of diameter at most 1, for convenience). Further, the natural counterpart of the assumption that equality is a congruence relation for the predicates and operations in a mathematical structure is the requirement that the predicates and operations in a metric structure be uniformly continuous with respect to the metric d. In the $[0,1]$-valued continuous setting, connectives are continuous functions on $[0,1]$ and quantifiers are sup and inf.

The analogy between this continuous version of first-order logic (CFO) for metric structures and the usual first-order logic (FOL) for ordinary structures is far reaching. In suitably phrased forms, CFO satisfies the compactness theorem, Löwenheim-Skolem theorems, diagram arguments, existence of saturated and homogeneous models, characterizations of quantifier elimination, Beth's definability theorem, the omitting types theorem, fundamental results of stability theory, and appropriate analogues of essentially all results in basic model theory of first-order logic. Moreover, CFO extends FOL: indeed, each mathematical structure treated in FOL can be viewed as a metric structure by taking the underlying metric d to be discrete ($d(a,b) = 1$ for distinct a, b). All these basic results true of CFO are thus framed as generalizations of the corresponding results for FOL.

A second type of justification for focusing on this continuous logic comes from its connection to applications of model theory in analysis and geometry. These often depend on an ultraproduct construction [11, 15] or, equivalently, the nonstandard hull construction (see [25, 24] and their references). This construction is widely used in functional analysis and also arises in metric space geometry (see [19], for example). The logic of positive bounded formulas was introduced in order to provide a model theoretic framework for the use of this ultraproduct (see [24]), which it does successfully. The continuous logic for metric structures that is presented here provides an equivalent background for this ultraproduct construction and it is easier to use. Writing positive bounded formulas to express statements from analysis and geometry is difficult and often feels unnatural; this goes much more smoothly in CFO. Indeed, continuous first-order logic provides model theorists and analysts with a common language; this is due to its being closely parallel to first-order logic while also using familiar constructs from analysis (*e.g.*, sup and inf in place of $\forall$ and $\exists$).

The purpose of this article is to present the syntax and semantics of this continuous logic for metric structures, to indicate some of its key theoretical features, and to show a few of its recent application areas.

In Sections 1 through 10 we develop the syntax and semantics of continuous logic for metric structures and present its basic properties. We have tried to make this material accessible without requiring any background beyond basic undergraduate mathematics. Sections 11 and 12 discuss imaginaries and omitting types; here our presentation is somewhat more brisk and full understanding may require some prior experience with model theory. Sections 13 and 14 sketch a treatment of quantifier elimination and stability, which are needed for the applications topics later in the paper; here we omit many proofs and depend on other articles for the details. Sections 15 through 18 indicate a few areas of mathematics to which continuous logic for metric structures has already been applied; these are taken from probability theory and functional analysis, and some background in these areas is expected of the reader.

The development of continuous logic for metric structures is very much a work in progress, and there are many open problems deserving of attention. What is presented in this article reflects work done over approximately the last three years in a series of collaborations among the authors. The material presented here was taught in two graduate topics courses offered during that time: a Fall 2004 course taught in Madison by Itaï Ben Yaacov and a Spring 2005 course taught in Urbana by Ward Henson. The authors are grateful to the students in those courses for their attention and help. The authors' research was partially supported by NSF Grants: Ben Yaacov, DMS-0500172; Berenstein and Henson, DMS-0100979 and DMS-0140677; Henson, DMS-0555904.

2 Metric structures and signatures

Let (M, d) be a complete, bounded metric space[1]. A *predicate* on M is a uniformly continuous function from M^n (for some $n \geq 1$) into some bounded interval in $\mathbb{R}$. A *function* or *operation* on M is a uniformly continuous function from M^n (for some $n \geq 1$) into M. In each case n is called the *arity* of the predicate or function.

A *metric structure* $\mathcal{M}$ based on (M, d) consists of a family $(R_i \mid i \in I)$ of predicates on M, a family $(F_j \mid j \in J)$ of functions on M, and a family $(a_k \mid k \in K)$ of distinguished elements of M. When we introduce such a metric structure, we will often denote it as

$$\mathcal{M} = (M, R_i, F_j, a_k \mid i \in I, j \in J, k \in K).$$

Any of the index sets I, J, K is allowed to be empty. Indeed, they might all be empty, in which case $\mathcal{M}$ is a pure bounded metric space.

The key restrictions on metric structures are: the metric space is *complete* and *bounded*, each predicate takes its values in a *bounded interval* of reals, and the functions and predicates are *uniformly continuous*. All of these restrictions play a role in making the theory work smoothly.

Our theory also applies to *many-sorted* metric structures, and they will appear as examples. However, in this article we will not explicitly bring them into our definitions and theorems, in order to avoid distracting notation.

2.1 Examples. We give a number of examples of metric structures to indicate the wide range of possibilities.

(1) A complete, bounded metric space (M, d) with no additional structure.
(2) A structure $\mathcal{M}$ in the usual sense from first-order logic. One puts the discrete metric on the underlying set ($d(a, b) = 1$ when a, b are distinct) and a relation is considered as a predicate taking values ("truth" values) in the set $\{0, 1\}$. So, in this sense the theory developed here is a generalization of first-order model theory.
(3) If (M, d) is an unbounded complete metric space with a distinguished element a, we may view (M, d) as a many-sorted metric structure $\mathcal{M}$; for example, we could take a typical sort to be a closed ball B_n of radius n around a, equipped with the metric obtained by restricting d. The inclusion mappings $I_{mn} \colon B_m \to B_n$

1 See the appendix to this section for some relevant basic facts about metric spaces.

$(m < n)$ should be functions in $\mathcal{M}$, in order to tie together the different sorts.

(4) The unit ball B of a Banach space X over $\mathbb{R}$ or $\mathbb{C}$: as functions we may take the maps $f_{\alpha\beta}$, defined by $f_{\alpha\beta}(x, y) = \alpha x + \beta y$, for each pair of scalars satisfying $|\alpha| + |\beta| \leq 1$; the norm may be included as a predicate, and we may include the additive identity 0_X as a distinguished element. Equivalently, X can be viewed as a many-sorted structure, with a sort for each ball of positive integer radius centered at 0, as indicated in the previous paragraph.

(5) Banach lattices: this is the result of expanding the metric structure corresponding to X as a Banach space (see the previous paragraph) by adding functions such as the absolute value operation on B as well as the positive and negative part operations. In section 17 of this article we discuss the model theory of some specific Banach lattices (namely, the L^p-spaces).

(6) Banach algebras: multiplication is included as an operation; if the algebra has a multiplicative identity, it may be included as a constant.

(7) C^*-algebras: multiplication and the *-map are included as operations.

(8) Hilbert spaces with inner product may be treated like the Banach space examples above, with the addition that the inner product is included as a binary predicate. (See section 15.)

(9) If $(\Omega, \mathcal{B}, \mu)$ is a probability space, we may construct a metric structure $\mathcal{M}$ from it, based on the metric space (M, d) in which M is the measure algebra of $(\Omega, \mathcal{B}, \mu)$ (elements of $\mathcal{B}$ modulo sets of measure 0) and d is defined to be the measure of the symmetric difference. As operations on M we take the Boolean operations $\cup, \cap, ^c$, as a predicate we take the measure μ, and as distinguished elements the 0 and 1 of M. In section 16 of this article we discuss the model theory of these metric structures.

Signatures

To each metric structure $\mathcal{M}$ we associate a *signature* L as follows. To each predicate R of $\mathcal{M}$ we associate a *predicate symbol* P and an integer $a(P)$ which is the arity of R; we denote R by $P^{\mathcal{M}}$. To each function F of $\mathcal{M}$ we associate a *function symbol* f and an integer $a(f)$ which is the arity of F; we denote F by $f^{\mathcal{M}}$. Finally, to each distinguished element a of $\mathcal{M}$ we associate a *constant symbol* c; we denote a by $c^{\mathcal{M}}$.

So, a signature L gives sets of predicate, function, and constant symbols, and associates to each predicate and function symbol its arity. In that respect, L is identical to a signature of first-order model theory. In addition, a signature for metric structures must specify more: for each predicate symbol P, it must provide a closed bounded interval I_P of real numbers and a modulus of uniform continuity[2] Δ_P. These should satisfy the requirements that $P^{\mathcal{M}}$ takes its values in I_P and that Δ_P is a modulus of uniform continuity for $P^{\mathcal{M}}$. In addition, for each function symbol f, L must provide a modulus of uniform continuity Δ_f, and this must satisfy the requirement that Δ_f is a modulus of uniform continuity for $f^{\mathcal{M}}$. Finally, L must provide a non-negative real number D_L which is a bound on the diameter of the complete metric space (M, d) on which $\mathcal{M}$ is based.[3] We sometimes denote the metric d given by $\mathcal{M}$ as $d^{\mathcal{M}}$; this would be consistent with our notation for the interpretation in $\mathcal{M}$ of the nonlogical symbols of L. However, we also find it convenient often to use the same notation "d" for the logical symbol representing the metric as well as for its interpretation in $\mathcal{M}$; this is consistent with usual mathematical practice and with the handling of the symbol $=$ in first-order logic.

When these requirements are all met and when the predicate, function, and constant symbols of L correspond exactly to the predicates, functions, and distinguished elements of which $\mathcal{M}$ consists, then we say $\mathcal{M}$ is an *L-structure*.

The key added features of a signature L in the metric structure setting are that L specifies (1) a bound on the diameter of the underlying metric space, (2) a modulus of uniform continuity for each predicate and function, and (3) a closed bounded interval of possible values for each predicate.

For simplicity, and without losing any generality, we will usually assume that our signatures L satisfy $D_L = 1$ and $I_P = [0, 1]$ for every predicate symbol P.

2.2 Remark. If $\mathcal{M}$ is an L-structure and A is a given closed subset of M^n, then $\mathcal{M}$ can be expanded by adding the predicate $x \mapsto \operatorname{dist}(x, A)$, where x ranges over M^n and dist denotes the distance function with respect to the maximum metric on the product space M^n. Note that only in very special circumstances may A itself be added to $\mathcal{M}$ as a predicate (in the form of the characteristic function χ_A of A); this could

2 See the appendix to this section for a discussion of this notion.
3 If L is many-sorted, each sort will have its own diameter bound.

be done only if χ_A were uniformly continuous, which forces A to be a positive distance from its complement in M^n.

Basic concepts such as *embedding* and *isomorphism* have natural definitions for metric structures:

2.3 Definition. Let L be a signature for metric structures and suppose $\mathcal{M}$ and $\mathcal{N}$ are L-structures.

An *embedding* from $\mathcal{M}$ into $\mathcal{N}$ is a metric space isometry

$$T\colon (M, d^{\mathcal{M}}) \to (N, d^{\mathcal{N}})$$

that commutes with the interpretations of the function and predicate symbols of L in the following sense:

Whenever f is an n-ary function symbol of L and $a_1, \ldots, a_n \in M$, we have

$$f^{\mathcal{N}}(T(a_1), \ldots, T(a_n)) = T(f^{\mathcal{M}}(a_1, \ldots, a_n));$$

whenever c is a constant symbol c of L, we have

$$c^{\mathcal{N}} = T(c^{\mathcal{M}});$$

and whenever P is an n-ary predicate symbol of L and $a_1, \ldots, a_n \in M$, we have

$$P^{\mathcal{N}}(T(a_1), \ldots, T(a_n)) = P^{\mathcal{M}}(a_1, \ldots, a_n).$$

An *isomorphism* is a surjective embedding. We say that $\mathcal{M}$ and $\mathcal{N}$ are *isomorphic*, and write $\mathcal{M} \cong \mathcal{N}$, if there exists an isomorphism between $\mathcal{M}$ and $\mathcal{N}$. (Sometimes we say *isometric isomorphism* to emphasize that isomorphisms must be distance preserving.) An *automorphism* of $\mathcal{M}$ is an isomorphism between $\mathcal{M}$ and itself.

$\mathcal{M}$ is a *substructure* of $\mathcal{N}$ (and we write $\mathcal{M} \subseteq \mathcal{N}$) if $M \subseteq N$ and the inclusion map from M into N is an embedding of $\mathcal{M}$ into $\mathcal{N}$.

Appendix

In this appendix we record some basic definitions and facts about metric spaces and uniformly continuous functions; they will be needed when we develop the semantics of continuous first-order logic. Proofs of the results we state here are straightforward and will mostly be omitted.

Let (M, d) be a metric space. We say this space is *bounded* if there is a real number B such that $d(x, y) \leq B$ for all $x, y \in M$. The *diameter* of (M, d) is the smallest such number B.

Suppose (M_i, d_i) are metric spaces for $i = 1, \dots, n$ and we take M to be the product $M = M_1 \times \cdots \times M_n$. In this article we will always regard M as being equipped with the maximum metric, defined for $x = (x_1, \dots, x_n), y = (y_1, \dots, y_n)$ by $d(x, y) = \max\{d_i(x_i, y_i) | i = 1, \dots, n\}$.

A *modulus of uniform continuity* is any function $\Delta\colon (0, 1] \to (0, 1]$.

If (M, d) and (M', d') are metric spaces and $f\colon M \to M'$ is any function, we say that $\Delta\colon (0, 1] \to (0, 1]$ is a *modulus of uniform continuity for* f if for every $\epsilon \in (0, 1]$ and every $x, y \in M$ we have

$$\text{(UC)} \qquad d(x, y) < \Delta(\epsilon) \quad \Longrightarrow \quad d'(f(x), f(y)) \le \epsilon.$$

We say f is *uniformly continuous* if it has a modulus of uniform continuity.

The precise way (UC) is stated makes the property Δ *is a modulus of uniform continuity for* f a topologically robust notion. For example, if $f\colon M \to M'$ is continuous and (UC) holds for a dense set of pairs (x, y), then it holds for all (x, y). In particular, if Δ is a modulus of uniform continuity for $f\colon M \to M'$ and we extend f in the usual way to a continuous function $\bar{f}\colon \overline{M} \to \overline{M'}$ (where $\overline{M}, \overline{M'}$ are completions of M, M', resp.), then, with this definition, Δ is a modulus of uniform continuity for the extended function $\bar{f}$.

If Δ is a function from $(0, \infty)$ to $(0, \infty)$ and it satisfies (UC) for all $\epsilon \in (0, \infty)$ and all $x, y \in M$, then we will often refer to Δ as a "modulus of uniform continuity" for f. In that case, f is uniformly continuous and the restriction of the function $\min(\Delta(\epsilon), 1)$ to $\epsilon \in (0, 1]$ is a modulus of uniform continuity according to the strict meaning we have chosen to assign to this phrase, so no confusion should result.

2.4 Proposition. *Suppose $f\colon M \to M'$ and $f'\colon M' \to M''$ are functions between metrics spaces M, M', M''. Suppose Δ is a modulus of uniform continuity for f and Δ' is a modulus of uniform continuity for f'. Then the composition $f' \circ f$ is uniformly continuous; indeed, for each $r \in (0, 1)$ the function $\Delta(r\Delta'(\epsilon))$ is a modulus of uniform continuity for $f' \circ f$.*

Let M, M' be metric spaces (with metrics d, d' resp.) and let f and $(f_n \mid n \ge 1)$ be functions from M into M'. Recall that $(f_n \mid n \ge 1)$ *converges uniformly to f on M* if

$$\forall \epsilon > 0\ \exists N\ \forall n > N\ \forall x \in M\ \big(d'(f_n(x), f(x)) \le \epsilon\big).$$

2.5 Proposition. *Let M, M', f and $(f_n \mid n \geq 1)$ be as above, and suppose $(f_n \mid n \geq 1)$ converges uniformly to f on M. If each of the functions $f_n \colon M \to M'$ is uniformly continuous, then f must also be uniformly continuous. Indeed, a modulus of uniform continuity for f can be obtained from moduli Δ_n for f_n, for each $n \geq 1$, and from a function $N \colon (0,1] \to \mathbb{N}$ that satisfies*

$$\forall \epsilon > 0 \; \forall n > N(\epsilon) \; \forall x \in M \left(d'(f_n(x), f(x)) \leq \epsilon \right).$$

Proof A modulus Δ for f may be defined as follows: given $\epsilon > 0$, take $\Delta(\epsilon) = \Delta_n(\epsilon/3)$ where $n = N(\epsilon/3) + 1$. □

2.6 Proposition. *Suppose $f, f_n \colon M \to M'$ and $f', f'_n \colon M' \to M''$ are functions $(n \geq 1)$ between metric spaces M, M', M''. If $(f_n \mid n \geq 1)$ converges uniformly to f on M and $(f'_n \mid n \geq 1)$ converges uniformly to f' on M', then $(f'_n \circ f_n \mid n \geq 1)$ converges uniformly to $f' \circ f$ on M.*

Fundamental to the continuous logic described in this article are the operations sup and inf on bounded sets of real numbers. We use these to define new functions from old, as follows. Suppose M, M' are metric spaces and $f \colon M \times M' \to \mathbb{R}$ is a bounded function. We define new functions $\sup_y f$ and $\inf_y f$ from M to $\mathbb{R}$ by

$$(\sup_y f)(x) = \sup\{f(x,y) | y \in M'\}$$

$$(\inf_y f)(x) = \inf\{f(x,y) | y \in M'\}$$

for all $x \in M$. Note that these new functions map M into the same closed bounded interval in $\mathbb{R}$ that contained the range of f. Our perspective is that $\sup_y$ and $\inf_y$ are quantifiers that bind or eliminate the variable y, analogous to the way $\forall$ and $\exists$ are used in ordinary first-order logic.

2.7 Proposition. *Suppose M, M' are metric spaces and f is a bounded uniformly continuous function from $M \times M'$ to $\mathbb{R}$. Let Δ be a modulus of uniform continuity for f. Then $\sup_y f$ and $\inf_y f$ are bounded uniformly continuous functions from M to $\mathbb{R}$, and Δ is a modulus of uniform continuity for both of them.*

Proof Fix $\epsilon > 0$ and consider $u, v \in M$ such that $d(u,v) < \Delta(\epsilon)$. Then for every $z \in M'$ we have

$$f(v,z) \leq f(u,z) + \epsilon \leq (\sup_y f)(u) + \epsilon.$$

Taking the sup over $z \in M'$ and interchanging the role of u and v yields

$$|(\sup_y f)(u) - (\sup_y f)(v)| \leq \epsilon.$$

The function $\inf_y f$ is handled similarly. □

2.8 Proposition. *Suppose M is a metric space and $f_s \colon M \to [0,1]$ is a uniformly continuous function for each s in the index set S. Let Δ be a common modulus of uniform continuity for $(f_s \mid s \in S)$. Then $\sup_s f_s$ and $\inf_s f_s$ are uniformly continuous functions from M to $[0,1]$, and Δ is a modulus of uniform continuity for both of them.*

Proof In the previous proof, take M' to be S with the discrete metric, and define $f(x,s) = f_s(x)$. □

2.9 Proposition. *Suppose M, M' are metric spaces and let $(f_n \mid n \geq 1)$ and f all be bounded functions from $M \times M'$ into $\mathbb{R}$. If $(f_n \mid n \geq 1)$ converges uniformly to f on $M \times M'$, then $(\sup_y f_n \mid n \geq 1)$ converges uniformly to $\sup_y f$ on M and $(\inf_y f_n \mid n \geq 1)$ converges uniformly to $\inf_y f$ on M.*

Proof Similar to the proof of Proposition 2.7. □

In many situations it is natural to construct a metric space as the quotient of a pseudometric space (M_0, d_0); here we mean that M_0 is a set and $d_0 \colon M_0 \times M_0 \to \mathbb{R}$ is a pseudometric. That is,

$$\begin{aligned} &d_0(x,x) = 0 \\ &d_0(x,y) = d_0(y,x) \geq 0 \\ &d_0(x,z) \leq d_0(x,y) + d_0(y,z) \end{aligned}$$

for all $x, y, z \in M_0$; these are the same conditions as in the definition of a metric, except that $d_0(x,y) = 0$ is allowed even when x, y are distinct.

If (M_0, d_0) is a pseudometric space, we may define an equivalence relation E on M_0 by $E(x,y) \Leftrightarrow d_0(x,y) = 0$. It follows from the triangle inequality that d_0 is E-invariant; that is, $d_0(x,y) = d_0(x',y')$ whenever xEx' and yEy'. Let M be the quotient set M_0/E and $\pi \colon M_0 \to M$ the quotient map, so $\pi(x)$ is the E-equivalence class of x, for each $x \in M_0$. Further, define d on M by setting $d(\pi(x), \pi(y)) = d_0(x,y)$ for any $x, y \in M_0$. Then (M,d) is a metric space and π is a distance preserving function from (M_0, d_0) onto (M,d). We will refer to (M,d) as the *quotient* metric space induced by (M_0, d_0).

Suppose (M_0, d_0) and (M_0', d_0') are pseudometric spaces with quotient metric spaces (M, d) and (M', d') and quotient maps π, π', respectively. Let $f_0\colon M_0 \to M_0'$ be any function. We say that f_0 is *uniformly continuous*, with modulus of uniform continuity Δ, if

$$d_0(x,y) < \Delta(\epsilon) \quad \Longrightarrow \quad d_0'(f_0(x), f_0(y)) \leq \epsilon$$

for all $x, y \in M_0$ and all $\epsilon \in (0,1]$. In that case it is clear that $d_0(x,y) = 0$ implies $d_0'(f_0(x), f_0(y)) = 0$ for all $x, y \in M_0$. Therefore we get a well defined quotient function $f\colon M \to M'$ by setting $f(\pi(x)) = \pi'(f_0(x))$ for all $x \in M_0$. Moreover, f is uniformly continuous with modulus of uniform continuity Δ.

The following results are useful in many places for expressing certain kinds of implications in continuous logic, and for reformulating the concept of uniform continuity.

2.10 Proposition. *Let $F, G\colon X \to [0,1]$ be arbitrary functions such that*

$$(\star) \qquad \forall \epsilon > 0\ \exists \delta > 0\ \forall x \in X\ \ (F(x) \leq \delta \Rightarrow G(x) \leq \epsilon)\,.$$

Then there exists an increasing, continuous function $\alpha\colon [0,1] \to [0,1]$ such that $\alpha(0) = 0$ and

$$(\star\star) \qquad \forall x \in X\ \ (G(x) \leq \alpha(F(x)))\,.$$

Proof Define a (possibly discontinuous) function $g\colon [0,1] \to [0,1]$ by

$$g(t) = \sup\{G(x) \mid F(x) \leq t\}$$

for $t \in [0,1]$. It is clear that g is increasing and that $G(x) \leq g(F(x))$ holds for all $x \in X$. Moreover, statement $(\star)$ implies that $g(0) = 0$ and that $g(t)$ converges to 0 as $t \to 0$.

To complete the proof we construct an increasing, continuous function $\alpha\colon [0,1] \to [0,1]$ such that $\alpha(0) = 0$ and $g(t) \leq \alpha(t)$ for all $t \in [0,1]$. Let $(t_n \mid n \in \mathbb{N})$ be any decreasing sequence in $[0,1]$ with $t_0 = 1$ and $\lim_{n\to\infty} t_n = 0$. Define $\alpha\colon [0,1] \to [0,1]$ by setting $\alpha(0) = 0$, $\alpha(1) = 1$, and $\alpha(t_n) = g(t_{n-1})$ for all $n \geq 1$, and by taking α to be linear on each interval of the form $[t_{n+1}, t_n]$, $n \in \mathbb{N}$. It is easy to check that α has the desired properties. For example, if $t_1 \leq t \leq t_0 = 1$ we have that $\alpha(t)$ is a convex combination of $g(1)$ and 1 so that $g(t) \leq g(1) \leq \alpha(t)$. Similarly, if $t_{n+1} \leq t \leq t_n$ and $n \geq 1$, we have that $\alpha(t)$ is a convex combination of $g(t_n)$ and $g(t_{n-1})$ so that $g(t) \leq g(t_n) \leq \alpha(t)$. Together with $g(0) = \alpha(0) = 0$, this shows that $g(t) \leq \alpha(t)$ for all $t \in [0, 1]$. □

2.11 Remark. Note that the converse to Proposition 2.10 is also true. Indeed, if statement ($\star\star$) holds and we fix $\epsilon > 0$, then taking

$$\delta = \sup\{s \mid \alpha(s) \leq \epsilon\} > 0$$

witnesses the truth of statement ($\star$).

2.12 Remark. The proof of Proposition 2.10 can be revised to show that the continuous function α can be chosen so that it only depends on the choice of an increasing function $\Delta\colon (0, 1] \to (0, 1]$ that witnesses the truth of statement ($\star$), in the sense that

$$\forall x \in X \ \ (F(x) \leq \Delta(\epsilon) \Rightarrow G(x) \leq \epsilon)$$

holds for each $\epsilon \in (0, 1]$. Given such a Δ, define $g\colon [0, 1] \to [0, 1]$ by $g(t) = \inf\{s \in (0, 1] \mid \Delta(s) > t\}$. It is easy to check that $g(0) = 0$ and that g is an increasing function. Moreover, for any $\epsilon > 0$ we have from the definition that $g(t) \leq \epsilon$ for any t in $[0, \Delta(\epsilon))$; therefore $g(t)$ converges to 0 as t tends to 0. Finally, we claim that $G(x) \leq g(F(x))$ holds for any $x \in X$. Otherwise we have $x \in X$ such that $g(F(x)) < G(x)$. The definition of g yields $s \in (0, 1]$ with $s < G(x)$ and $\Delta(s) > F(x)$; this contradicts our assumptions.

Now α is constructed from g as in the proof of Proposition 2.10. This yields an increasing, continuous function $\alpha\colon [0, 1] \to [0, 1]$ with $\alpha(0) = 0$ such that whenever $F, G\colon X \to [0, 1]$ are functions satisfying

$$\forall x \in X \ \ (F(x) \leq \Delta(\epsilon) \Rightarrow G(x) \leq \epsilon)$$

for each $\epsilon \in (0, 1]$, then we have

$$\forall x \in X \ \ (G(x) \leq \alpha(F(x)))\,.$$

3 Formulas and their interpretations

Fix a signature L for metric structures, as described in the previous section. As indicated there (see page 6), we assume for simplicity of notation that $D_L = 1$ and that $I_P = [0, 1]$ for every predicate symbol P.

Symbols of L

Among the *symbols* of L are the predicate, function, and constant symbols; these will be referred to as the *nonlogical* symbols of L and the remaining ones will be called the *logical* symbols of L. Among the logical symbols is a symbol d for the metric on the underlying metric space of an L-structure; this is treated formally as equivalent to a predicate symbol of arity 2. The logical symbols also include an infinite set V_L of *variables*; usually we take V_L to be countable, but there are situations in which it is useful to permit a larger number of variables. The remaining logical symbols consist of a symbol for each continuous function $u\colon [0,1]^n \to [0,1]$ of finitely many variables $n \geq 1$ (these play the role of connectives) and the symbols sup and inf, which play the role of quantifiers in this logic.

The *cardinality* of L, denoted $\mathrm{card}(L)$, is the smallest infinite cardinal number $\geq$ the number of nonlogical symbols of L.

Terms of L

Terms are formed inductively, exactly as in first-order logic. Each variable and constant symbol is an L-term. If f is an n-ary function symbol and $t_1, \dots, t_n$ are L-terms, then $f(t_1, \dots, t_n)$ is an L-term. All L-terms are constructed in this way.

Atomic formulas of L.

The *atomic formulas* of L are the expressions of the form $P(t_1, \dots, t_n)$, in which P is an n-ary predicate symbol of L and $t_1, \dots, t_n$ are L-terms; as well as $d(t_1, t_2)$, in which t_1 and t_2 are L-terms.

Note that the logical symbol d for the metric is treated formally as a binary predicate symbol, exactly analogous to how the equality symbol $=$ is treated in first-order logic.

Formulas of L

Formulas are also constructed inductively, and the basic structure of the induction is similar to the corresponding definition in first-order logic. Continuous functions play the role of connectives and sup and inf are used formally in the way that quantifiers are used in first-order logic. The precise definition is as follows:

3.1 Definition. The class of L-formulas is the smallest class of expressions satisfying the following requirements:

(1) Atomic formulas of L are L-formulas.
(2) If $u\colon [0,1]^n \to [0,1]$ is continuous and $\varphi_1, \dots, \varphi_n$ are L-formulas, then $u(\varphi_1, \dots, \varphi_n)$ is an L-formula.
(3) If φ is an L-formula and x is a variable, then $\sup_x \varphi$ and $\inf_x \varphi$ are L-formulas.

3.2 Remark. We have chosen to take all continuous functions on $[0,1]$ as our connectives. This is both too restrictive (see section 9 in which we want to close our set of formulas under certain kinds of limits, in order to develop a good notion of *definability*) and too general (see section 6). We made this choice in order to introduce formulas as early and as directly as possible.

An L-formula is *quantifier free* if it is generated inductively from atomic formulas without using the last clause; *i.e.*, neither $\sup_x$ nor $\inf_x$ are used.

Many syntactic notions from first-order logic can be carried over word for word into this setting. We will assume that this has been done by the reader for many such concepts, including *subformula* and *syntactic substitution* of a term for a variable, or a formula for a subformula, and so forth.

Free and *bound* occurrences of variables in L-formulas are defined in a way similar to how this is done in first-order logic. Namely, an occurrence of the variable x is bound if lies within a subformula of the form $\sup_x \varphi$ or $\inf_x \varphi$, and otherwise it is free.

An *L-sentence* is an L-formula that has no free variables.

When t is a term and the variables occurring in it are among the variables $x_1, \dots, x_n$ (which we always take to be distinct in this context), we indicate this by writing t as $t(x_1, \dots, x_n)$.

Similarly, we write an L-formula as $\varphi(x_1, \dots, x_n)$ to indicate that its free variables are among $x_1, \dots, x_n$.

Prestructures

It is common in mathematics to construct a metric space as the quotient of a pseudometric space or as the completion of such a quotient, and the same is true of metric structures. For that reason we need to consider what we will call *prestructures* and to develop the semantics of continuous logic for them.

As above, we take L to be a fixed signature for metric structures. Let (M_0, d_0) be a pseudometric space, satisfying the requirement that its diameter is $\leq D_L$. (That is, $d_0(x, y) \leq D_L$ for all $x, y \in M_0$.) An *L-prestructure* $\mathcal{M}_0$ based on (M_0, d_0) is a structure consisting of the following data:

(1) for each predicate symbol P of L (of arity n) a function $P^{\mathcal{M}_0}$ from M_0^n into I_P that has Δ_P as a modulus of uniform continuity;
(2) for each function symbol f of L (of arity n) a function $f^{\mathcal{M}_0}$ from M_0^n into M_0 that has Δ_f as a modulus of uniform continuity; and
(3) for each constant symbol c of L an element $c^{\mathcal{M}_0}$ of M_0.

Given an L-prestructure $\mathcal{M}_0$, we may form its *quotient* prestructure as follows. Let (M, d) be the quotient metric space induced by (M_0, d_0) with quotient map $\pi \colon M_0 \to M$. Then

(1) for each predicate symbol P of L (of arity n) define $P^{\mathcal{M}}$ from M^n into I_P by setting $P^{\mathcal{M}}(\pi(x_1), \dots, \pi(x_n)) = P^{\mathcal{M}_0}(x_1, \dots, x_n)$ for each $x_1, \dots, x_n \in M_0$;
(2) for each function symbol f of L (of arity n) define $f^{\mathcal{M}}$ from M^n into M by setting $f^{\mathcal{M}}(\pi(x_1), \dots, \pi(x_n)) = \pi(f^{\mathcal{M}_0}(x_1, \dots, x_n))$ for each $x_1, \dots, x_n \in M_0$;
(3) for each constant symbol c of L define $c^{\mathcal{M}} = \pi(c^{\mathcal{M}_0})$.

It is obvious that (M, d) has the same diameter as (M_0, d_0). Also, as noted in the appendix to section 2, for each predicate symbol P and each function symbol f of L, the predicate $P^{\mathcal{M}}$ is well defined and has Δ_P as a modulus of uniform continuity and the function $f^{\mathcal{M}}$ is well defined and has Δ_f as a modulus of uniform continuity. In other words, this defines an L-prestructure (which we will denote as $\mathcal{M}$) based on the (possibly not complete) metric space (M, d).

Finally, we may define an L-structure $\mathcal{N}$ by taking a *completion* of $\mathcal{M}$. This is based on a complete metric space (N, d) that is a completion of (M, d), and its additional structure is defined in the following natural way (made possible by the fact that the predicates and functions given by $\mathcal{M}$ are uniformly continuous):

(1) for each predicate symbol P of L (of arity n) define $P^{\mathcal{N}}$ from N^n into I_P to be the unique such function that extends $P^{\mathcal{M}}$ and is continuous;

(2) for each function symbol f of L (of arity n) define $f^{\mathcal{N}}$ from N^n into N to be the unique such function that extends $f^{\mathcal{M}}$ and is continuous;
(3) for each constant symbol c of L define $c^{\mathcal{N}} = c^{\mathcal{M}}$.

It is obvious that (N, d) has the same diameter as (M, d). Also, as noted in the appendix to section 2, for each predicate symbol P and each function symbol f of L, the predicate $P^{\mathcal{N}}$ has Δ_P as a modulus of uniform continuity and the function $f^{\mathcal{N}}$ has Δ_f as a modulus of uniform continuity. In other words, $\mathcal{N}$ is an L-structure.

Semantics

Let $\mathcal{M}$ be any L-prestructure, with $(M, d^{\mathcal{M}})$ as its underlying pseudometric space, and let A be a subset of M. We extend L to a signature $L(A)$ by adding a new constant symbol $c(a)$ to L for each element a of A. We extend the interpretation given by $\mathcal{M}$ in a canonical way, by taking the interpretation of $c(a)$ to be equal to a itself for each $a \in A$. We call $c(a)$ the *name* of a in $L(A)$. Indeed, we will often write a instead of $c(a)$ when no confusion can result from doing so.

Given an $L(M)$-term $t(x_1, \dots, x_n)$, we define, exactly as in first-order logic, the interpretation of t in $\mathcal{M}$, which is a function $t^{\mathcal{M}} \colon M^n \to M$.

We now come to the key definition in continuous logic for metric structures, in which the semantics of this logic is defined. For each $L(M)$-sentence σ, we define *the value of* σ *in* $\mathcal{M}$. This value is a real number in the interval $[0, 1]$ and it is denoted $\sigma^{\mathcal{M}}$. The definition is by induction on formulas. Note that in the definition all terms mentioned are $L(M)$-terms in which no variables occur.

3.3 Definition. (1) $\big(d(t_1, t_2)\big)^{\mathcal{M}} = d^{\mathcal{M}}(t_1^{\mathcal{M}}, t_2^{\mathcal{M}})$ for any t_1, t_2;
(2) for any n-ary predicate symbol P of L and any $t_1, \dots, t_n$,
$$\big(P(t_1, \dots, t_n)\big)^{\mathcal{M}} = P^{\mathcal{M}}(t_1^{\mathcal{M}}, \dots, t_n^{\mathcal{M}});$$
(3) for any $L(M)$-sentences $\sigma_1, \dots, \sigma_n$ and any continuous function $u \colon [0, 1]^n \to [0, 1]$,
$$\big(u(\sigma_1, \dots, \sigma_n)\big)^{\mathcal{M}} = u(\sigma_1^{\mathcal{M}}, \dots, \sigma_n^{\mathcal{M}});$$
(4) for any $L(M)$-formula $\varphi(x)$,
$$\big(\sup_x \varphi(x)\big)^{\mathcal{M}}$$

is the supremum in $[0,1]$ of the set $\{\varphi(a)^{\mathcal{M}} \mid a \in M\}$;

(5) for any $L(M)$-formula $\varphi(x)$,

$$\big(\inf_x \varphi(x)\big)^{\mathcal{M}}$$

is the infimum in $[0,1]$ of the set $\{\varphi(a)^{\mathcal{M}} \mid a \in M\}$.

3.4 Definition. Given an $L(M)$-formula $\varphi(x_1, \ldots, x_n)$ we let $\varphi^{\mathcal{M}}$ denote the function from M^n to $[0,1]$ defined by

$$\varphi^{\mathcal{M}}(a_1, \ldots, a_n) = \big(\varphi(a_1, \ldots, a_n)\big)^{\mathcal{M}}.$$

A key fact about formulas in continuous logic is that they define uniformly continuous functions. Indeed, the modulus of uniform continuity for the predicate does not depend on $\mathcal{M}$ but only on the data given by the signature L.

3.5 Theorem. *Let $t(x_1, \ldots, x_n)$ be an L-term and $\varphi(x_1, \ldots, x_n)$ an L-formula. Then there exist functions Δ_t and Δ_φ from $(0,1]$ to $(0,1]$ such that for any L-prestructure $\mathcal{M}$, Δ_t is a modulus of uniform continuity for the function $t^{\mathcal{M}} \colon M^n \to M$ and Δ_φ is a modulus of uniform continuity for the predicate $\varphi^{\mathcal{M}} \colon M^n \to [0,1]$.*

Proof The proof is by induction on terms and then induction on formulas. The basic tools concerning uniform continuity needed for the induction steps in the proof are given in the appendix to section 2. □

3.6 Remark. The previous result is the counterpart in this logic of the Perturbation Lemma in the logic of positive bounded formulas with the approximate semantics. See [24, Proposition 5.15].

3.7 Theorem. *Let $\mathcal{M}_0$ be an L-prestructure with underlying pseudometric space (M_0, d_0); let $\mathcal{M}$ be its quotient L-structure with quotient map $\pi \colon M_0 \to M$ and let $\mathcal{N}$ be the L-structure that results from completing $\mathcal{M}$ (as explained on page 15). Let $t(x_1, \ldots, x_n)$ be any L-term and $\varphi(x_1, \ldots, x_n)$ be any L-formula. Then:*

(1) $t^{\mathcal{M}}(\pi(a_1), \ldots, \pi(a_n)) = t^{\mathcal{M}_0}(a_1, \ldots, a_n)$ for all $a_1, \ldots, a_n \in M_0$;

(2) $t^{\mathcal{N}}(b_1, \ldots, b_n) = t^{\mathcal{M}}(b_1, \ldots, b_n)$ for all $b_1, \ldots, b_n \in M$.

(3) $\varphi^{\mathcal{M}}(\pi(a_1), \ldots, \pi(a_n)) = \varphi^{\mathcal{M}_0}(a_1, \ldots, a_n)$ for all $a_1, \ldots, a_n \in M_0$;

(4) $\varphi^{\mathcal{N}}(b_1, \ldots, b_n) = \varphi^{\mathcal{M}}(b_1, \ldots, b_n)$ for all $b_1, \ldots, b_n \in M$.

Proof The proofs are by induction on terms and then induction on formulas. In handling the quantifier cases in (3) the key is that the quotient map π is surjective. For the quantifier cases in (4), the key is that the functions $\varphi^{\mathcal{N}}$ are continuous and that $\mathcal{M}$ is dense in $\mathcal{N}$. □

3.8 Caution. Note that we only use words such as *structure* when the underlying metric space is *complete*. In some constructions this means that we must take a metric space completion at the end. Theorem 3.7 shows that this preserves all properties expressible in continuous logic.

Logical equivalence

3.9 Definition. Two L-formulas $\varphi(x_1, \ldots, x_n)$ and $\psi(x_1, \ldots, x_n)$ are said to be *logically equivalent* if

$$\varphi^{\mathcal{M}}(a_1, \ldots, a_n) = \psi^{\mathcal{M}}(a_1, \ldots, a_n)$$

for every L-structure $\mathcal{M}$ and every $a_1, \ldots, a_n \in M$.

If $\varphi(x_1, \ldots, x_n)$ and $\psi(x_1, \ldots, x_n)$ are L-formulas, we can extend the preceding definition by taking the *logical distance* between φ and ψ to be the supremum of all numbers

$$|\varphi^{\mathcal{M}}(a_1, \ldots, a_n) - \psi^{\mathcal{M}}(a_1, \ldots, a_n)|$$

where $\mathcal{M}$ is any L-structure and $a_1, \ldots, a_n \in M$. This defines a pseudo-metric on the set of all formulas with free variables among $x_1, \ldots, x_n$, and two formulas are logically equivalent if and only if the logical distance between them is 0.

3.10 Remark. Note that by Theorem 3.7, we could use L-prestructures in place of L-structures in the preceding Definition without changing the meaning of the concepts defined.

3.11 Remark. (Size of the space of L-formulas)
Some readers may be concerned that the set of L-formulas is be too large, because we allow all continuous functions as connectives. What matters, however, is the size of a set of L-formulas that is dense in the set of all L-formulas with respect to the logical distance defined in the previous paragraph. By Weierstrass's Theorem, there is a countable set of functions from $[0, 1]^n$ to $[0, 1]$ that is dense in the set of all continuous functions, with respect to the sup-distance between such functions. (The

sup-distance between f and g is the supremum of $|f(x) - g(x)|$ as x ranges over the common domain of f, g.) If we only use such connectives in building L-formulas, then (a) the total number of formulas that are constructed is card(L), and (b) any L-formula can be approximated arbitrarily closely in logical distance by a formula constructed using the restricted connectives. Thus the density character of the set of L-formulas with respect to logical distance is always $\leq$ card(L). (We explore this topic in more detail in section 6.)

Conditions of L

An *L-condition* E is a formal expression of the form $\varphi = 0$, where φ is an L-formula. We call E *closed* if φ is a sentence. If $x_1, \dots, x_n$ are distinct variables, we write an L-condition as $E(x_1, \dots, x_n)$ to indicate that it has the form $\varphi(x_1, \dots, x_n) = 0$ (in other words, that the free variables of E are among $x_1, \dots, x_n$).

If E is the $L(M)$-condition $\varphi(x_1, \dots, x_n) = 0$ and $a_1, \dots, a_n$ are in M, we say E is *true of* $a_1, \dots, a_n$ *in* $\mathcal{M}$ and write $\mathcal{M} \models E[a_1, \dots, a_n]$ if $\varphi^{\mathcal{M}}(a_1, \dots, a_n) = 0$.

3.12 Definition. Let E_i be the L-condition $\varphi_i(x_1, \dots, x_n) = 0$, for $i = 1, 2$. We say that E_1 and E_2 are *logically equivalent* if for every L-structure $\mathcal{M}$ and every $a_1, \dots, a_n$ we have

$$\mathcal{M} \models E_1[a_1, \dots, a_n] \quad \text{iff} \quad \mathcal{M} \models E_2[a_1, \dots, a_n].$$

3.13 Remark. When φ and ψ are formulas, it is convenient to introduce the expression $\varphi = \psi$ as an abbreviation for the condition $|\varphi - \psi| = 0$. (Note that $u\colon [0,1]^2 \to [0,1]$ defined by $u(t_1, t_2) = |t_1 - t_2|$ is a connective.) Since each real number $r \in [0,1]$ is a connective (thought of as a constant function), expressions of the form $\varphi = r$ will thereby be regarded as conditions for any L-formula φ and $r \in [0,1]$. Note that the interpretation of $\varphi = \psi$ is semantically correct; namely for any L-structure $\mathcal{M}$ and elements a of M, $|\varphi - \psi|^{\mathcal{M}}(a) = 0$ if and only if $\varphi^{\mathcal{M}}(a) = \psi^{\mathcal{M}}(a)$.

Similarly, we introduce the expressions $\varphi \leq \psi$ and $\psi \geq \varphi$ as abbreviations for certain conditions. Let $\dot{-}\colon [0,1]^2 \to [0,1]$ be the connective defined by $\dot{-}(t_1, t_2) = \max(t_1 - t_2, 0) = t_1 - t_2$ if $t_1 \geq t_2$ and 0 otherwise. Usually we write $t_1 \dot{-} t_2$ in place of $\dot{-}(t_1, t_2)$. We take $\varphi \leq \psi$ and $\psi \geq \varphi$ to be abbreviations for the condition $\varphi \dot{-} \psi = 0$. (See section 6, where this connective plays a central role.) In $[0,1]$-valued logic, the condition $\varphi \leq \psi$ can be seen as family of implications, from the condition $\psi \leq r$ to the condition $\varphi \leq r$ for each $r \in [0,1]$.

4 Model theoretic concepts

Fix a signature L for metric structures. In this section we introduce several of the most fundamental model theoretic concepts and discuss some of their basic properties.

4.1 Definition. A *theory* in L is a set of closed L-conditions. If T is a theory in L and $\mathcal{M}$ is an L-structure, we say that $\mathcal{M}$ *is a model* of T and write $\mathcal{M} \models T$ if $\mathcal{M} \models E$ for every condition E in T. We write $\mathrm{Mod}_L(T)$ for the collection of all L-structures that are models of T. (If L is clear from the context, we write simply $\mathrm{Mod}(T)$.)

If $\mathcal{M}$ is an L-structure, the *theory of* $\mathcal{M}$, denoted $\mathrm{Th}(\mathcal{M})$, is the set of closed L-conditions that are true in $\mathcal{M}$. If T is a theory of this form, it will be called *complete*.

If T is an L-theory and E is a closed L-condition, we say E *is a logical consequence of* T and write $T \models E$ if $\mathcal{M} \models E$ holds for every model $\mathcal{M}$ of T.

4.2 Caution. Note that we only use words such as *model* when the underlying metric space is *complete*. Theorem 3.7 shows that whenever T is an L-theory and $\mathcal{M}_0$ is an L-prestructure such that $\varphi^{\mathcal{M}_0} = 0$ for every condition $\varphi = 0$ in T, then the completion of the canonical quotient of $\mathcal{M}_0$ is indeed a *model* of T.

4.3 Definition. Suppose that $\mathcal{M}$ and $\mathcal{N}$ are L-structures.

(1) We say that $\mathcal{M}$ and $\mathcal{N}$ are *elementarily equivalent*, and write $\mathcal{M} \equiv \mathcal{N}$, if $\sigma^{\mathcal{M}} = \sigma^{\mathcal{N}}$ for all L-sentences σ. Equivalently, this holds if $\mathrm{Th}(\mathcal{M}) = \mathrm{Th}(\mathcal{N})$.

(2) If $\mathcal{M} \subseteq \mathcal{N}$ we say that $\mathcal{M}$ is an *elementary substructure of* $\mathcal{N}$, and write $\mathcal{M} \preceq \mathcal{N}$, if whenever $\varphi(x_1, \ldots, x_n)$ is an L-formula and $a_1, \ldots, a_n$ are elements of M, we have

$$\varphi^{\mathcal{M}}(a_1, \ldots, a_n) = \varphi^{\mathcal{N}}(a_1, \ldots, a_n).$$

In this case, we also say that $\mathcal{N}$ is an *elementary extension* of $\mathcal{M}$.

(3) A function F from a subset of M into N is an *elementary map* from $\mathcal{M}$ into $\mathcal{N}$ if whenever $\varphi(x_1, \ldots, x_n)$ is an L-formula and $a_1, \ldots, a_n$ are elements of the domain of F, we have

$$\varphi^{\mathcal{M}}(a_1, \ldots, a_n) = \varphi^{\mathcal{N}}(F(a_1), \ldots, F(a_n)).$$

(4) An *elementary embedding* of $\mathcal{M}$ into $\mathcal{N}$ is a function from all of M into N that is an elementary map from $\mathcal{M}$ into $\mathcal{N}$.

4.4 Remark. (1) Every elementary map from one metric structure into another is distance preserving.
(2) The collection of elementary maps is closed under composition and formation of the inverse.
(3) Every isomorphism between metric structures is an elementary embedding.

In the following result we refer to sets $\mathcal{S}$ of L-formulas that are *dense with respect to logical distance.* (See page 18.) That is, such a set has the following property: for any L-formula $\varphi(x_1, \ldots, x_n)$ and any $\epsilon > 0$ there is $\psi(x_1, \ldots, x_n)$ in $\mathcal{S}$ such that for any L-structure $\mathcal{M}$ and any $a_1, \ldots, a_n \in M$

$$|\varphi^{\mathcal{M}}(a_1, \ldots, a_n) - \psi^{\mathcal{M}}(a_1, \ldots, a_n)| \leq \epsilon.$$

4.5 Proposition. *(Tarski-Vaught Test for $\preceq$) Let $\mathcal{S}$ be any set of L-formulas that is dense with respect to logical distance. Suppose $\mathcal{M}$, $\mathcal{N}$ are L-structures with $\mathcal{M} \subseteq \mathcal{N}$. The following statements are equivalent:*

(1) $\mathcal{M} \preceq \mathcal{N}$;
(2) *For every L-formula $\varphi(x_1, \ldots, x_n, y)$ in $\mathcal{S}$ and $a_1, \ldots, a_n \in M$,*

$$\inf\{\varphi^{\mathcal{N}}(a_1, \ldots, a_n, b) \mid b \in N\} = \inf\{\varphi^{\mathcal{N}}(a_1, \ldots, a_n, c) \mid c \in M\}$$

Proof If (1) holds, then we may conclude (2) for the set of all L-formulas directly from the meaning of $\preceq$. Indeed, if $\varphi(x_1, \ldots, x_n, y)$ is any L-formula and $a_1, \ldots, a_n \in A$, then from (1) we have

$$\inf\{\varphi^{\mathcal{N}}(a_1, \ldots, a_n, b) \mid b \in N\} = \big(\inf_y \varphi(a_1, \ldots, a_n, y)\big)^{\mathcal{N}} =$$

$$\big(\inf_y \varphi(a_1, \ldots, a_n, y)\big)^{\mathcal{M}} = \inf\{\varphi^{\mathcal{M}}(a_1, \ldots, a_n, c) \mid c \in M\} =$$

$$\inf\{\varphi^{\mathcal{N}}(a_1, \ldots, a_n, c) \mid c \in M\}.$$

For the converse, suppose (2) holds for a set $\mathcal{S}$ that is dense in the set of all L-formulas with respect to logical distance. First we will prove that (2) holds for the set of all L-formulas. Let $\varphi(x_1, \dots, x_n, y)$ be any L-formula. Given $\epsilon > 0$, let $\psi(x_1, \dots, x_n, y)$ be an element of $\mathcal{S}$ that approximates $\varphi(x_1, \dots, x_n, y)$ to within ϵ in logical distance. Let $a_1, \dots, a_n$ be elements of M. Then we have

$$\begin{aligned}
&\inf\{\varphi^{\mathcal{N}}(a_1, \dots, a_n, b) \mid b \in N\} \\
&\leq \inf\{\psi^{\mathcal{N}}(a_1, \dots, a_n, b) \mid b \in N\} + \epsilon \\
&= \inf\{\psi^{\mathcal{N}}(a_1, \dots, a_n, c) \mid c \in M\} + \epsilon \\
&\leq \inf\{\varphi^{\mathcal{N}}(a_1, \dots, a_n, c) \mid c \in M\} + 2\epsilon.
\end{aligned}$$

Letting ϵ tend to 0 and recalling $M \subseteq N$ we obtain the desired equality for $\varphi(x_1, \dots, x_n, y)$.

Now assume that (2) holds for the set of all L-formulas. One proves the equivalence

$$\psi^{\mathcal{M}}(a_1, \dots, a_n) = \psi^{\mathcal{N}}(a_1, \dots, a_n)$$

(for all $a_1, \dots, a_n$ in M) by induction on the complexity of ψ, using (2) to cover the case when ψ begins with sup or inf. □

5 Ultraproducts and compactness

First we discuss ultrafilter limits in topology. Let X be a topological space and let $(x_i)_{i \in I}$ be a family of elements of X. If D is an ultrafilter on I and $x \in X$, we write

$$\lim_{i,D} x_i = x$$

and say *x is the D-limit of $(x_i)_{i \in I}$* if for every neighborhood U of x, the set $\{ i \in I \mid x_i \in U \}$ is in the ultrafilter D. A basic fact from general topology is that X is a compact Hausdorff space if and only if for every family $(x_i)_{i \in I}$ in X and every ultrafilter D on I the D-limit of $(x_i)_{i \in I}$ exists and is unique.

The following lemmas are needed below when we connect ultrafilter limits and the semantics of continuous logic.

5.1 Lemma. *Suppose X, X' are topological spaces and $F\colon X \to X'$ is continuous. For any family $(x_i)_{i \in I}$ from X and any ultrafilter D on I, we have that*

$$\lim_{i,D} x_i = x \quad \Longrightarrow \quad \lim_{i,D} F(x_i) = F(x)$$

where the ultrafilter limits are taken in X and X' respectively.

Proof Let U be an open neighborhood of $F(x)$ in X'. Since F is continuous, $F^{-1}(U)$ is open in X, and it contains x. If x is the D-limit of $(x_i)_{i\in I}$, there exists $A \in D$ such that for all $i \in A$ we have $x_i \in F^{-1}(U)$ and hence $F(x_i) \in U$. □

5.2 Lemma. *Let X be a closed, bounded interval in $\mathbb{R}$. Let S be any set and let $(F_i \mid i \in I)$ be a family of functions from S into X. Then, for any ultrafilter D on I*

$$\sup_x \big(\lim_{i,D} F_i(x)\big) \quad \le \quad \lim_{i,D} \big(\sup_x F_i(x)\big), \text{ and}$$

$$\inf_x \big(\lim_{i,D} F_i(x)\big) \quad \ge \quad \lim_{i,D} \big(\inf_x F_i(x)\big).$$

where in both cases, $\sup_x$ *and* $\inf_x$ *are taken over $x \in S$. Moreover, for each $\epsilon > 0$ there exist $(x_i)_{i\in I}$ and $(y_i)_{i\in I}$ in S such that*

$$\lim_{i,D} F_i(x_i) + \epsilon \quad \ge \quad \lim_{i,D} \big(\sup_x F_i(x)\big), \text{ and}$$

$$\lim_{i,D} F_i(y_i) - \epsilon \quad \le \quad \lim_{i,D} \big(\inf_x F_i(x)\big).$$

Proof We prove the statements involving sup; the inf statements are proved similarly (or by replacing each F_i by its negative).

Let $r_i = \sup_x F_i(x)$ for each $i \in I$ and let $r = \lim_{i,D} r_i$. For each $\epsilon > 0$, let $A(\epsilon) \in D$ be such that $r - \epsilon < r_i < r + \epsilon$ for every $i \in A(\epsilon)$.

First we show $\sup_x \lim_{i,D} F_i(x) \le r$. For each $i \in A(\epsilon)$ and $x \in S$ we have $F_i(x) \le r_i < r + \epsilon$. Hence the D-limit of $(F_i(x))_{i\in I}$ is $\le r + \epsilon$. Letting ϵ tend to 0 gives the desired inequality.

For the other sup statement, fix $\epsilon > 0$ and for each $i \in I$ choose $x_i \in S$ so that $r_i \le F_i(x_i) + \epsilon/2$. Then for $i \in A(\epsilon/2)$ we have $r \le F_i(x_i) + \epsilon$. Taking the D-limit gives the desired inequality. □

Ultraproducts of metric spaces

Let $((M_i, d_i) \mid i \in I)$ be a family of bounded metric spaces, all having diameter $\le K$. Let D be an ultrafilter on I. Define a function d on the Cartesian product $\prod_{i\in I} M_i$ by

$$d(x, y) = \lim_{i,D} d_i(x_i, y_i)$$

where $x = (x_i)_{i\in I}$ and $y = (y_i)_{i\in I}$. This D-limit is taken in the interval $[0, K]$. It is easy to check that d is a pseudometric on $\prod_{i\in I} M_i$.

For $x, y \in \prod_{i\in I} M_i$, define $x \sim_D y$ to mean that $d(x, y) = 0$. Then $\sim_D$ is an equivalence relation, so we may define

$$\Big(\prod_{i\in I} M_i\Big)_D = \Big(\prod_{i\in I} M_i\Big)\Big/{\sim_D}.$$

The pseudometric d on $\prod_{i\in I} M_i$ induces a metric on this quotient space, and we also denote this metric by d.

The space $(\prod_{i\in I} M_i)_D$ with the induced metric d is called the *D-ultraproduct of* $((M_i, d_i) \mid i \in I)$. We denote the equivalence class of $(x_i)_{i\in I} \in \prod_{i\in I} M_i$ under $\sim_D$ by $((x_i)_{i\in I})_D$.

If $(M_i, d_i) = (M, d)$ for every $i \in I$, the space $(\prod_{i\in I} M_i)_D$ is called the *D-ultrapower of* M and it is denoted $(M)_D$. In this situation, the map $T\colon M \to (M)_D$ defined by $T(x) = ((x_i)_{i\in I})_D$, where $x_i = x$ for every $i \in I$, is an isometric embedding. It is called the *diagonal embedding* of M into $(M)_D$.

A particular case of importance is the D-ultrapower of a compact metric space (M, d). In that case the diagonal embedding of M into $(M)_D$ is surjective. Indeed, if $(x_i)_{i\in I} \in M^I$ and x is the D-limit of the family $(x_i)_{i\in I}$, which exists since (M, d) is compact, then it is easy to show that $((x_i)_{i\in I})_D = T(x)$. In particular, any ultrapower of a closed bounded interval may be canonically identified with the interval itself.

Since we require that structures are based on *complete* metric spaces, it is useful to note that every ultraproduct of such spaces is itself complete.

5.3 Proposition. *Let $((M_i, d_i) \mid i \in I)$ be a family of complete, bounded metric spaces, all having diameter $\leq K$. Let D be an ultrafilter on I and let (M, d) be the D-ultraproduct of $((M_i, d_i) \mid i \in I)$. The metric space (M, d) is complete.*

Proof Let $(x^k)_{k\geq 1}$ be a Cauchy sequence in (M, d). Without loss of generality we may assume that $d(x^k, x^{k+1}) < 2^{-k}$ holds for all $k \geq 1$; that is, to prove (M, d) complete it suffices to show that all such Cauchy sequences have a limit. For each $k \geq 1$ let x^k be represented by the family $(x_i^k)_{i\in I}$. For each $m \geq 1$ let A_m be the set of all $i \in I$ such that $d_i(x_i^k, x_i^{k+1}) < 2^{-k}$ holds for all $k = 1, \dots, m$. Then the sets $(A_m)_{m\geq 1}$ form a decreasing chain and all of them are in D.

We define a family $(y_i)_{i\in I}$ that will represent the limit of the sequence

$(x^k)_{k\geq 1}$ in (M,d). If $i \notin A_1$, then we take y_i to be an arbitrary element of M_i. If for some $m \geq 1$ we have $i \in A_m \setminus A_{m+1}$, then we set $y_i = x_i^{m+1}$. If $i \in A_m$ holds for all $m \geq 1$, then $(x_i^m)_{m\geq 1}$ is a Cauchy sequence in the complete metric space (M_i, d_i) and we take y_i to be its limit.

An easy calculation shows that for each $m \geq 1$ and each $i \in A_m$ we have $d_i(x_i^m, y_i) \leq 2^{-m+1}$. It follows that $((y_i)_{i\in I})_D$ is the limit in the ultraproduct (M,d) of the sequence $(x^k)_{k\geq 1}$. □

Ultraproducts of functions

Suppose $((M_i,d_i) \mid i \in I)$ and $((M_i',d_i') \mid i \in I)$ are families of metric spaces, all of diameter $\leq K$. Fix $n \geq 1$ and suppose $f_i \colon M_i^n \to M_i'$ is a uniformly continuous function for each $i \in I$. Moreover, suppose the single function $\Delta \colon (0,1] \to (0,1]$ is a modulus of uniform continuity for all of the functions f_i. Given an ultrafilter D on I, we define a function

$$\Big(\prod_{i\in I} f_i\Big)_D \colon \Big(\prod_{i\in I} M_i\Big)_D^n \to \Big(\prod_{i\in I} M_i'\Big)_D$$

as follows. If for each $k = 1,\dots,n$ we have $(x_i^k)_{i\in I} \in \prod_{i\in I} M_i$, we define

$$\Big(\prod_{i\in I} f_i\Big)_D \big(((x_i^1)_{i\in I})_D, \dots, ((x_i^n)_{i\in I})_D\big) = \Big(\Big(f_i(x_i^1,\dots,x_i^n)\Big)_{i\in I}\Big)_D.$$

We claim that this defines a uniformly continuous function that also has Δ as its modulus of uniform continuity. For simplicity of notation, suppose $n = 1$. Fix $\epsilon > 0$. Suppose the distance between $\big((x_i)_{i\in I}\big)_D$ and $\big((y_i)_{i\in I}\big)_D$ in the ultraproduct $\Big(\prod_{i\in I} M_i\Big)_D$ is $< \Delta(\epsilon)$. There must exist $A \in D$ such that for all $i \in A$ we have $d_i(x_i,y_i) < \Delta(\epsilon)$. Since Δ is a modulus of uniform continuity for all of the functions f_i, it follows that $d_i'(f_i(x_i), f_i(y_i)) \leq \epsilon$ for all $i \in A$. Hence the distance in the ultraproduct $\Big(\prod_{i\in I} M_i'\Big)_D$ between $\big((f(x_i))_{i\in I}\big)_D$ and $\big((f(y_i))_{i\in I}\big)_D$ must be $\leq \epsilon$. This shows that $\Big(\prod_{i\in I} f_i\Big)_D$ is well defined and that it has Δ as a modulus of uniform continuity. (Note that the precise form of our definition of "modulus of uniform continuity" played a role in this argument.)

Ultraproducts of L-structures

Let $(\mathcal{M}_i \mid i \in I)$ be a family of L-structures and let D be an ultrafilter on I. Suppose the underlying metric space of $\mathcal{M}_i$ is (M_i, d_i). Since there is a uniform bound on the diameters of these metric spaces, we may form their D-ultraproduct. For each function symbol f of L, the functions $f^{\mathcal{M}_i}$ all have the same modulus of uniform continuity Δ_f. Therefore the D-ultraproduct of this family of functions is well defined. The same is true if we consider a predicate symbol P of L. Moreover, the functions $P^{\mathcal{M}_i}$ all have their values in $[0,1]$, whose D-ultrapower can be identified with $[0,1]$ itself; thus the D-ultraproduct of $(P^{\mathcal{M}_i} \mid i \in I)$ can be regarded as a $[0,1]$-valued function on M.

Therefore we may define the *D-ultraproduct of the family* $(\mathcal{M}_i \mid i \in I)$ *of L-structures* to be the L-structure $\mathcal{M}$ that is specified as follows:

The underlying metric space of $\mathcal{M}$ is given by the ultraproduct of metric spaces

$$M = \Big(\prod_{i\in I} M_i\Big)_D.$$

For each predicate symbol P of L, the interpretation of P in $\mathcal{M}$ is given by the ultraproduct of functions

$$P^{\mathcal{M}} = \Big(\prod_{i\in I} P^{\mathcal{M}_i}\Big)_D$$

which maps M^n to $[0,1]$. For each function symbol f of L, the interpretation of f in $\mathcal{M}$ is given by the ultraproduct of functions

$$f^{\mathcal{M}} = \Big(\prod_{i\in I} f^{\mathcal{M}_i}\Big)_D$$

which maps M^n to M. For each constant symbol c of L, the interpretation of c in $\mathcal{M}$ is given by

$$c^{\mathcal{M}} = \big((c^{\mathcal{M}_i})_{i\in I}\big)_D.$$

The discussion above shows that this defines $\mathcal{M}$ to be a well-defined L-structure. We call $\mathcal{M}$ the *D-ultraproduct of the family* $(\mathcal{M}_i \mid i \in I)$ and denote it by

$$\mathcal{M} = \Big(\prod_{i\in I} \mathcal{M}_i\Big)_D.$$

If all of the L-structures $\mathcal{M}_i$ are equal to the same structure $\mathcal{M}_0$, then $\mathcal{M}$ is called the *D-ultrapower of* $\mathcal{M}_0$ and is denoted by

$$\big(\mathcal{M}_0\big)_D.$$

This ultraproduct construction finds many applications in functional analysis (see [24] and its references) and in metric space geometry (see

[19]). Its usefulness is partly explained by the following theorem, which is the analogue in this setting of the well known result in first-order logic proved by J. Łos. This is sometimes known as the *Fundamental Theorem of Ultraproducts.*

5.4 Theorem. *Let $(\mathcal{M}_i \mid i \in I)$ be a family of L-structures. Let D be any ultrafilter on I and let $\mathcal{M}$ be the D-ultraproduct of $(\mathcal{M}_i \mid i \in I)$. Let $\varphi(x_1, \dots, x_n)$ be an L-formula. If $a_k = \big((a_i^k)_{i\in I}\big)_D$ are elements of $\mathcal{M}$ for $k = 1, \dots, n$, then*

$$\varphi^{\mathcal{M}}(a_1, \dots, a_n) = \lim_{i,D} \varphi^{\mathcal{M}_i}(a_i^1, \dots, a_i^n).$$

Proof The proof is by induction on the complexity of φ. Basic facts about ultrafilter limits (discussed at beginning of this section) are used in the proof. □

5.5 Corollary. *If $\mathcal{M}$ is an L-structure and $T\colon M \to (M)_D$ is the diagonal embedding, then T is an elementary embedding of $\mathcal{M}$ into $(\mathcal{M})_D$.*

Proof From Theorem 5.4. □

5.6 Corollary. *If $\mathcal{M}$ and $\mathcal{N}$ are L-structures and they have isomorphic ultrapowers, then $\mathcal{M} \equiv \mathcal{N}$.*

Proof Immediate from the preceding result. □

The converse of the preceding corollary is also true, in a strong form:

5.7 Theorem. *If $\mathcal{M}$ and $\mathcal{N}$ are L-structures and $\mathcal{M} \equiv \mathcal{N}$, then there exists an ultrafilter D such that $(\mathcal{M})_D$ is isomorphic to $(\mathcal{N})_D$.*

The preceding result is an extension of the Keisler-Shelah Theorem from ordinary model theory. (See [36] and Chapter 6 in [13].) A detailed proof of the analogous result for normed space structures and the approximate logic of positive bounded formulas is given in [24, Chapter 10], and that argument can be readily adapted to continuous logic for metric structures.

There are characterizations of elementary equivalence that are slightly more complex to state than Theorem 5.7 but are much easier to prove, such as the following: $\mathcal{M} \equiv \mathcal{N}$ if and only if $\mathcal{M}$ and $\mathcal{N}$ have isomorphic elementary extensions that are each constructed as the union of an infinite sequence of successive ultrapowers. Theorem 5.7 has the positive feature that it connects continuous logic in a direct way to application areas in which the ultrapower construction is important, such as the theory of Banach spaces and other areas of functional analysis. Indeed, the result shows that the mathematical properties a metric structure and its ultrapowers share in common are exactly those that can be expressed by sentences of the continuous analogue of first-order logic. Theorem 5.7 also yields a characterization of axiomatizable classes of metric structures (5.14 below) whose statement is simpler than would otherwise be the case.

Compactness theorem

5.8 Theorem. *Let T be an L-theory and $\mathcal{C}$ a class of L-structures. Assume that T is finitely satisfiable in $\mathcal{C}$. Then there exists an ultraproduct of structures from $\mathcal{C}$ that is a model of T.*

Proof Let Λ be the set of finite subsets of T. Let $\lambda \in \Lambda$, and write $\lambda = \{E_1, \dots, E_n\}$. By assumption there is an L-structure $\mathcal{M}_\lambda$ in $\mathcal{C}$ such that $\mathcal{M}_\lambda \models E_j$ for all $j = 1, \dots, n$.

For each $E \in T$, let $S(E)$ be the set of all $\lambda \in \Lambda$ such that $E \in \lambda$. Note that the collection of sets $\{S(E) \mid E \in T\}$ has the finite intersection property. Hence there is an ultrafilter D on Λ that contains this collection.

Let

$$\mathcal{M} \;=\; \Big(\prod_{\lambda\in\Lambda} \mathcal{M}_\lambda\Big)_D.$$

Note that if $\lambda \in S(E)$, then $\mathcal{M}_\lambda \models E$. It follows from Theorem 5.4 that $\mathcal{M} \models E$ for every $E \in T$. In other words, the ultraproduct $\mathcal{M}$ of structures from $\mathcal{C}$ is a model of T. □

In many applications it is useful to note that the Compactness Theorem remains true even if the finite satisfiability hypothesis is weakened to an approximate version. This is an immediate consequence of basic properties of the semantics for continuous logic.

5.9 Definition. For any set Σ of L-conditions, Σ^+ is the set of all conditions $\varphi \leq 1/n$ such that $\varphi = 0$ is an element of Σ and $n \geq 1$.

5.10 Corollary. *Let T be an L-theory and $\mathcal{C}$ a class of L-structures. Assume that T^+ is finitely satisfiable in $\mathcal{C}$. Then there exists an ultraproduct of structures from $\mathcal{C}$ that is a model of T.*

Proof This follows immediately from Theorem 5.8, because T and T^+ obviously have the same models. □

The next result is a version of the Compactness Theorem for formulas. In it we allow an arbitrary family $(x_j \mid j \in J)$ of possible free variables.

5.11 Definition. Let T be an L-theory and $\Sigma(x_j \mid j \in J)$ a set of L-conditions. We say that Σ is *consistent with* T if for every finite subset F of Σ there exists a model $\mathcal{M}$ of T and elements a of M such that for every condition E in F we have $\mathcal{M} \models E[a]$. (Here a is a finite tuple suitable for the free variables in members of F.)

5.12 Corollary. *Let T be an L-theory and $\Sigma(x_j \mid j \in J)$ a set of L-conditions, and assume that Σ^+ is consistent with T. Then there is a model $\mathcal{M}$ of T and elements $(a_j \mid j \in J)$ of M such that*

$$\mathcal{M} \models E[a_j \mid j \in J]$$

for every L-condition E in Σ.

Proof Let $(c_j \mid j \in J)$ be new constants and consider the signature $L(\{c_j \mid j \in J\})$. This corollary is proved by applying the Compactness Theorem to the set $T \cup \Sigma^+(c_j \mid j \in J)$ of closed $L(\{c_j \mid j \in J\})$-conditions. As noted in the proof of the previous result, anything satisfying Σ^+ will also satisfy Σ. □

Axiomatizability of classes of structures

5.13 Definition. Suppose that $\mathcal{C}$ is a class of L-structures. We say that $\mathcal{C}$ is *axiomatizable* if there exists a set T of closed L-conditions such that $\mathcal{C} = \mathrm{Mod}_L(T)$. When this holds for T, we say that T is a set of *axioms* for $\mathcal{C}$ in L.

In this section we characterize axiomatizability in continuous logic using ultraproducts. The ideas are patterned after a well known characterization of axiomatizability in first-order logic due to Keisler [31]. (See Corollary 6.1.16 in [13].)

5.14 Proposition. *Suppose that $\mathcal{C}$ is a class of L-structures. The following statements are equivalent:*

(1) *$\mathcal{C}$ is axiomatizable in L;*
(2) *$\mathcal{C}$ is closed under isomorphisms and ultraproducts, and its complement, $\{\mathcal{M} \mid \mathcal{M}$ is an L-structure not in $\mathcal{C}\}$, is closed under ultrapowers.*

Proof (1)⇒(2) follows from the Fundamental Theorem of Ultraproducts.

To prove (2)⇒(1), we let T be the set of closed L-conditions that are satisfied by every structure in $\mathcal{C}$. We claim that T is a set of axioms for $\mathcal{C}$. To prove this, suppose $\mathcal{M}$ is an L-structure such that $\mathcal{M} \models T$.

We claim that $\mathrm{Th}(\mathcal{M})^+$ is finitely satisfiable in $\mathcal{C}$. If not, there exist L-sentences $\sigma_1, \dots, \sigma_n$ and $\epsilon > 0$ such that $\sigma_j^{\mathcal{M}} = 0$ for all $j = 1, \dots, n$, but such that for any $\mathcal{N} \in \mathcal{C}$, we have $\sigma_j^{\mathcal{N}} \geq \epsilon$ for some $j = 1, \dots, n$. This means that the condition $\max(\sigma_1, \dots, \sigma_n) \geq \epsilon$ is in T but is not satisfied in $\mathcal{M}$, which is a contradiction.

So $\mathrm{Th}(\mathcal{M})^+$ is finitely satisfiable in $\mathcal{C}$. By the Compactness Theorem this yields an ultraproduct $\mathcal{M}'$ of structures from $\mathcal{C}$ such that M' is a model of $\mathrm{Th}(\mathcal{M})^+$. One sees easily that this implies $\mathcal{M}' \equiv \mathcal{M}$. Theorem 5.7, the extension of the Keisler-Shelah theorem to this continuous logic, yields an ultrafilter D such that $(\mathcal{M}')_D$ and $(\mathcal{M})_D$ are isomorphic. Statement (2) implies that $\mathcal{M}$ is in $\mathcal{C}$. □

5.15 Remark. The proof of Proposition 5.14 contains the following useful elementary result: let $\mathcal{C}$ be a class of L-structures and let T be the set of all closed L-conditions E such that $\mathcal{M} \models E$ holds for all $\mathcal{M} \in \mathcal{C}$. Then, every model of T is elementarily equivalent to some ultraproduct of structures from $\mathcal{C}$.

6 Connectives

Recall that in our definition of *formulas* for continuous logic, we took a *connective* to be a continuous function from $[0,1]^n$ to $[0,1]$, for some $n \geq 1$. This choice is somewhat arbitrary; from one point of view it is too general, and from another it is too restrictive. We begin to discuss these issues in this section.

Here we continue to limit ourselves to *finitary* connectives, and our intention is to limit the connectives we use when building formulas. We consider a restricted set of formulas to be adequate if every formula can be "uniformly approximated" by formulas from the restricted set.

6.1 Definition. A system of connectives is a family $\mathcal{F} = (F_n \mid n \geq 1)$ where each F_n is a set of connectives $f\colon [0,1]^n \to [0,1]$. We say that $\mathcal{F}$ is *closed* if it is closed under arbitrary substitutions; more precisely:

(1) For each n, F_n contains the projection $\pi^n_j\colon [0,1]^n \to [0,1]$ onto the j^{th} coordinate for each $j = 1,\dots,n$.
(2) For each n and m, if $u \in F_n$, and $v_1,\dots,v_n \in F_m$, then the function $w\colon [0,1]^m \to [0,1]$ defined by $w(t) = u(v_1(t),\dots,v_n(t))$, where t denotes an element of $[0,1]^m$, is in F_m.

Note that each system $\mathcal{F}$ of connectives generates a smallest closed system of connectives $\overline{\mathcal{F}}$. We say that $\mathcal{F}$ is *full* if $\overline{\mathcal{F}}$ is uniformly dense in the system of all connectives; that is, for any $\epsilon > 0$ and any connective $f(t_1,\dots,t_n)$, there is a connective $g \in \overline{F}_n$ such that

$$|f(t_1,\dots,t_n) - g(t_1,\dots,t_n)| \leq \epsilon$$

for all $(t_1,\dots,t_n) \in [0,1]^n$.

6.2 Definition. Given a system $\mathcal{F}$ of connectives, we define the collection of $\mathcal{F}$*-restricted* formulas by induction:

(1) Atomic formulas are $\mathcal{F}$-restricted formulas.
(2) If $u \in F_n$ and $\varphi_1,\dots,\varphi_n$ are $\mathcal{F}$-restricted formulas, then $u(\varphi_1,\dots,\varphi_n)$ is an $\mathcal{F}$-restricted formula.
(3) If φ is an $\mathcal{F}$-restricted formula, so are $\sup_x \varphi$ and $\inf_x \varphi$.

The importance of full sets of connectives is that the restricted formulas made using them are dense in the set of all formulas, with respect to the logical distance between L-formulas, as the next result states. (See page 18.)

6.3 Theorem. *Assume that $\mathcal{F}$ is a full system of connectives. Then for any $\epsilon > 0$ and any L-formula $\varphi(x_1,\dots,x_n)$, there is an $\mathcal{F}$-restricted L-formula $\psi(x_1,\dots,x_n)$ such that for all L-structures $\mathcal{M}$ one has*

$$|\varphi^{\mathcal{M}}(a_1,\dots,a_n) - \psi^{\mathcal{M}}(a_1,\dots,a_n)| \leq \epsilon$$

for all $a_1,\dots,a_n \in M$.

Proof By induction on formulas. □

Next we show that there is a very simple system of connectives that is full. In particular this system is countable, so there are at most card(L) many restricted L-formulas for this system of connectives. It follows from the previous result that the collection of all L formulas has density at most card(L) with respect to the logical distance between L-formulas.

6.4 Definition. We define a binary function $\dot{-}\colon \mathbb{R}^{\geq 0} \times \mathbb{R}^{\geq 0} \to \mathbb{R}^{\geq 0}$ by:

$$x \dot{-} y = \begin{cases} (x-y) & \text{if } x \geq y \\ 0 & \text{otherwise.} \end{cases}$$

Note that if $x, y \in [0,1]$, then $x \dot{-} y \in [0,1]$, so the restriction of $\dot{-}$ is a connective; we use $\dot{-}$ to denote this connective as well as its extension to all of $\mathbb{R}^{\geq 0}$.

There are many well known identities involving $\dot{-}$. For example, note that $x \dot{-} y = (x+z) \dot{-} (y+z)$ and $((x \dot{-} y) \dot{-} z) = x \dot{-} (y+z)$ for all $x, y, z \in \mathbb{R}^{\geq 0}$.

6.5 Definition. Let $\mathcal{F}_0 = (F_n \mid n \geq 1)$ where $F_1 = \{0, 1, x/2\}$ (with $0, 1$ treated as constant functions of one variable), $F_2 = \{\dot{-}\}$, and all other F_n are empty. We call the connectives in $\mathcal{F}_0$ *restricted*.

The following connectives belong to $\overline{F}_2$:

$$\min(t_1, t_2) = t_1 \dot{-} (t_1 \dot{-} t_2)$$

$$\max(t_1, t_2) = 1 \dot{-} (\min(1 \dot{-} t_1, 1 \dot{-} t_2))$$

$$|t_1 - t_2| = \max(t_1 \dot{-} t_2, t_2 \dot{-} t_1)$$

$$\min(t_1 + t_2, 1) = 1 \dot{-} ((1 \dot{-} t_1) \dot{-} t_2)$$

$$t_1 \dot{-} (mt_2) = \underbrace{((\dots(t_1 \dot{-} t_2) \dot{-} \dots) \dot{-} t_2)}_{m\ times}.$$

Every dyadic fraction $m2^{-n}$ in $[0,1]$ is an element of $\overline{F}_1$.

6.6 Proposition. *The system of connectives $\mathcal{F}_0$ is full.*

Proof Let D be the set of dyadic fractions $m2^{-n}$ in $[0,1]$. First we show that for each distinct x, y in D, the set $\{(g(x), g(y)) \mid g \in \overline{F}_1\}$ includes all pairs (a,b) in D^2. Fix numbers $x, y, a, b \in D$ with $x < y$ and $a \geq b$.

Choose $m \in \mathbb{N}$ such that $a < m(y - x)$. Now let $g\colon [0,1] \to [0,1]$ be defined by

$$g(t) = \max(a \dot{-} m(t \dot{-} x), b).$$

Then $g \in \overline{F}_1$, and we also see that $g(x) = a$ and $g(y) = b$. If $a < b$ we can achieve the same result by using $1 \dot{-} a$ and $1 \dot{-} b$ in place of a, b in the construction of g, and then using the function $1 \dot{-} g(t)$.

Now we prove density by arguing exactly as in the proof of the lattice version of the Stone-Weierstrass Theorem on $[0,1]^n$. (See [17, pages 241–242].) □

6.7 Notation. By a *restricted formula* we mean an $\mathcal{F}_0$-restricted formula, where $\mathcal{F}_0$ is the system of connectives in Definition 6.5.

6.8 Definition. A formula is in *prenex form* if it is of the form

$$Q^1_{x_1} Q^2_{x_2} \dots Q^n_{x_n} \psi$$

where ψ is a quantifier free formula and each Q^i is either sup or inf.

6.9 Proposition. *Every restricted formula is equivalent to a restricted formula in prenex form.*

Proof By induction on formulas. The proof proceeds like the usual proof in first-order logic. The main point of the proof is that the connective $\dot{-}$ is monotone in its arguments, increasing in the first and decreasing in the second. □

Existential conditions

Since the family of restricted formulas is uniformly dense in the family of all formulas, it follows from the previous result that every L-formula can be uniformly approximated by formulas in prenex form. This gives us a way of introducing analogues of the usual syntactic classes into continuous logic. For example, an L-formula is defined to be an inf-*formula* if it is approximated arbitrarily closely by prenex formulas of the form $\inf_{x_1} \dots \inf_{x_n} \psi$, where ψ is quantifier free. A condition $\varphi = 0$ is defined to be *existential* if φ is an inf-formula. Other syntactic forms for L-conditions are defined similarly.

7 Constructions of models

Unions of chains

If Λ is a linearly ordered set, a Λ-*chain* of L-structures is a family of L-structures $(\mathcal{M}_\lambda \mid \lambda \in \Lambda)$ such that $\mathcal{M}_\lambda \subseteq \mathcal{M}_\eta$ for $\lambda < \eta$. If this holds, we can define the union of $(\mathcal{M}_\lambda \mid \lambda \in \Lambda)$ as an L-prestructure in an obvious way. (Note that for each function symbol or predicate symbol S, all of the interpretations $S^{\mathcal{M}_\lambda}$ have the same modulus of uniform continuity Δ_S, guaranteeing that the union of $(S^{\mathcal{M}_\lambda} \mid \lambda \in \Lambda)$ will also have Δ_S as a modulus of uniform continuity.) This union is based on a metric space, but it may not be complete. After taking the completion we get an L-structure that we will refer to as the *union* of the chain and that we will denote by $\bigcup_{\lambda\in\Lambda} \mathcal{M}_\lambda$.

Caution: In general, if the ordered set Λ has countable cofinality, the set-theoretic union $\bigcup_{\lambda\in\Lambda} M_\lambda$ will be a dense *proper* subset of the underlying set of the L-structure $\bigcup_{\lambda\in\Lambda} \mathcal{M}_\lambda$.

7.1 Definition. A chain of structures $(\mathcal{M}_\lambda \mid \lambda \in \Lambda)$ is called an *elementary chain* if $\mathcal{M}_\lambda \preceq \mathcal{M}_\eta$ for all $\lambda < \eta$.

7.2 Proposition. *If* $(\mathcal{M}_\lambda \mid \lambda \in \Lambda)$ *is an elementary chain and* $\lambda \in \Lambda$, *then* $\mathcal{M}_\lambda \preceq \bigcup_{\lambda\in\Lambda} \mathcal{M}_\lambda$.

Proof Use the Tarski-Vaught test (Proposition 4.5). □

Löwenheim-Skolem theorem

Recall that the *density character* of a topological space is the smallest cardinality of a dense subset of the space. For example, a space is separable if and only if its density character is $\leq \aleph_0$. If A is a topological space, we denote its density character by density(A).

7.3 Proposition. *(Downward Löwenheim-Skolem Theorem)*
Let κ *be an infinite cardinal number and assume* card$(L) \leq \kappa$. *Let* $\mathcal{M}$ *be an* L*-structure and suppose* $A \subseteq M$ *has* density$(A) \leq \kappa$. *Then there exists a substructure* $\mathcal{N}$ *of* $\mathcal{M}$ *such that*

(1) $\mathcal{N} \preceq \mathcal{M}$;
(2) $A \subseteq N \subseteq M$;
(3) density$(N) \leq \kappa$.

Proof Let A_0 be a dense subset of A of cardinality at most κ. By suitably enlarging A_0, we may obtain a prestructure N_0 such that $A_0 \subseteq N_0 \subseteq M$ and $\mathrm{card}(N_0) \leq \kappa$ and such that the following closure property also holds: for every restricted L-formula $\varphi(x_1, \ldots, x_n, x_{n+1})$ and every rational $\epsilon > 0$, if $\varphi^{\mathcal{M}}(a_1, \ldots, a_n, c) \leq \epsilon$ with $a_k \in N_0$ for $k = 1, \ldots, n$ and $c \in M$, then there exists $b \in N_0$ such that $\varphi^{\mathcal{M}}(a_1, \ldots, a_n, b) \leq \epsilon$. It is possible to do this while maintaining the claimed cardinality bounds because L has at most κ many restricted formulas.

Let N be the closure of N_0 in M. By considering atomic formulas in the closure property above, one shows that there is a substructure $\mathcal{N}$ of $\mathcal{M}$ that is based on N. Continuity of formulas and uniform density of restricted formulas show that $\mathcal{N} \subseteq \mathcal{M}$ satisfy the Tarski-Vaught test (Proposition 4.5) and hence that $\mathcal{N} \preceq \mathcal{M}$. □

Saturated structures

7.4 Definition. Let $\Gamma(x_1, \ldots, x_n)$ be a set of L-conditions and let $\mathcal{M}$ be an L-structure. We say that $\Gamma(x_1, \ldots, x_n)$ is *satisfiable in* $\mathcal{M}$ if there exist elements $a_1, \ldots, a_n$ of $\mathcal{M}$ such that $\mathcal{M} \models \Gamma[a_1, \ldots, a_n]$.

7.5 Definition. Let $\mathcal{M}$ be an L-structure and let κ be an infinite cardinal. We say that $\mathcal{M}$ is *κ-saturated* if the following statement holds: whenever $A \subseteq M$ has cardinality $< \kappa$ and $\Gamma(x_1, \ldots, x_n)$ is a set of $L(A)$-conditions, if every finite subset of Γ is satisfiable in $(\mathcal{M}, a)_{a \in A}$, then the entire set Γ is satisfiable in $(\mathcal{M}, a)_{a \in A}$.

It is straightforward using ultraproducts to prove the existence of ω_1-saturated L-structures when L is countable, and we will do that next. This is the only degree of saturation that one needs for many applications.

Recall that an ultrafilter D is said to be *countably incomplete* if D is not closed under countable intersections. It is equivalent to require the existence of elements J_n of D for each $n \in \mathbb{N}$ such that the intersection $\bigcap_{n \in \mathbb{N}} J_n$ is the empty set. Evidently any non-principal ultrafilter on a countable set is countably incomplete.

7.6 Proposition. *Let L be a signature with* $\mathrm{card}(L) = \omega$ *and let D be a countably incomplete ultrafilter on a set Λ. Then for every family $(\mathcal{M}_\lambda \mid \lambda \in \Lambda)$ of L-structures, $(\prod_{\lambda \in \Lambda} \mathcal{M}_\lambda)_D$ is ω_1-saturated.*

Proof In order to simplify the notation, we verify that $(\prod_{\lambda\in\Lambda}\mathcal{M}_\lambda)_D$ satisfies the statement in Definition 7.5 for $n = 1$.

We have to prove the following statement: if $A \subseteq (\prod_{\lambda\in\Lambda}\mathcal{M}_\lambda)_D$ is countable and $\Gamma(x)$ is a set of $L(A)$-formulas such that every finite subset of $\Gamma(x)$ is satisfiable in $\big((\prod_{\lambda\in\Lambda}\mathcal{M}_\lambda)_D,\, a\big)_{a\in A}$, then the entire set $\Gamma(x)$ is satisfied in $\big((\prod_{\lambda\in\Lambda}\mathcal{M}_\lambda)_D,\, a\big)_{a\in A}$. For each $a \in A$, let

$$u(a) = (u_\lambda(a) \mid \lambda \in \Lambda) \in \prod_{\lambda\in\Lambda}\mathcal{M}_\lambda$$

be such that $a = \big((u_\lambda(a))_{\lambda\in\Lambda}\big)_D$. Note that

$$\Big(\ (\prod_{\lambda\in\Lambda}\mathcal{M}_\lambda)_D,\ a\ \Big)_{a\in A} = \Big(\ \prod_{\lambda\in\Lambda}(\mathcal{M}_\lambda, u_\lambda(a))_{a\in A}\ \Big)_D.$$

Thus, since L is an arbitrary countable signature and A is also countable, it suffices to prove the following simpler statement:

If $\Gamma(x)$ is a set of L-conditions and every finite subset of $\Gamma(x)$ is satisfiable in $(\prod_{\lambda\in\Lambda}\mathcal{M}_\lambda)_D$, then $\Gamma(x)$ is satisfiable in $(\prod_{\lambda\in\Lambda}\mathcal{M}_\lambda)_D$.

So, suppose every finite subset of $\Gamma(x)$ is satisfiable in $(\prod_{\lambda\in\Lambda}\mathcal{M}_\lambda)_D$. Since L is countable, we may write

$$\Gamma(x) = \{\, \varphi_n(x) = 0 \mid n \in \mathbb{N}\,\}.$$

Since D is countably incomplete, we can fix a descending chain of elements of D

$$\Lambda = \Lambda_0 \supseteq \Lambda_1 \supseteq \dots$$

such that $\bigcap_{k\in\mathbb{N}}\Lambda_k = \emptyset$.

Let $X_0 = \Lambda$ and for each positive integer k define

$$X_k = \Lambda_k \cap \ \Big\{\ \lambda \in \Lambda \ \mid\ \mathcal{M}_\lambda \models \inf_x \max(\varphi_1, \dots, \varphi_k) \le \frac{1}{k+1}\ \Big\}.$$

Then $X_k \in D$ by Theorem 5.4, since

$$(\prod_{\lambda\in\Lambda}\mathcal{M}_\lambda)_D \models \inf_x \max(\varphi_1, \dots, \varphi_k) = 0.$$

We then have $X_k \supseteq X_{k+1}$ for every $k \in \mathbb{N}$ and $\bigcap_{k\in\mathbb{N}} X_k = \emptyset$, so for each $\lambda \in \Lambda$ there exists a largest positive integer $k(\lambda)$ such that $\lambda \in X_{k(\lambda)}$.

We now define an element $a = (a(\lambda))_{\lambda\in\Lambda} \in \prod_{\lambda\in\Lambda}\mathcal{M}_\lambda$ such that

$$\Big(\prod_{\lambda\in\Lambda}\mathcal{M}_\lambda\Big)_D \models \Gamma[(a)_D].$$

For $\lambda \in \Lambda$, if $k(\lambda) = 0$, let $a(\lambda)$ be any element in M_λ; otherwise, let $a(\lambda)$ be such that

$$\mathcal{M}_\lambda \models \left\{\max(\varphi_1, \dots, \varphi_{k(\lambda)}) \leq \frac{1}{k(\lambda)}\right\}[a(\lambda)].$$

If $k \in \mathbb{N}$ and $\lambda \in X_k$, we have $k \leq k(\lambda)$, so $\mathcal{M}_\lambda \models \left(\varphi_k \leq \frac{1}{k(\lambda)}\right)[a(\lambda)]$. It follows from Theorem 5.4 that $(\prod_{\lambda \in \Lambda} \mathcal{M}_\lambda)_D \models \Gamma[(a)_D]$. □

In saturated structures the meaning of L-conditions can be analyzed using the usual quantifiers $\forall$ and $\exists$, as the next result shows:

7.7 Proposition. *Let $\mathcal{M}$ be an L-structure and suppose $E(x_1, \dots, x_m)$ is the L-condition*

$$\left(Q^1_{y_1} \dots Q^n_{y_n} \varphi(x_1, \dots, x_m, y_1, \dots, y_n)\right) = 0$$

where each Q^i is either inf *or* sup *and φ is quantifier free.*

Let $\mathcal{E}(x_1, \dots, x_m)$ be the mathematical statement

$$\widetilde{Q}^1 y_1 \dots \widetilde{Q}^n y_n \left(\varphi(x_1, \dots, x_m, y_1, \dots, y_n) = 0\right)$$

where each $\widetilde{Q}^i y_i$ is $\exists y_i$ if $Q^i_{y_i}$ is $\inf_{y_i}$ and is $\forall y_i$ if $Q^i_{y_i}$ is $\sup_{y_i}$.

If $\mathcal{M}$ is ω-saturated, then for any elements $a_1, \dots, a_m$ of M, we have

$$\mathcal{M} \models E[a_1, \dots, a_m] \textit{ if and only if } \mathcal{E}(a_1, \dots, a_n) \textit{ is true in } \mathcal{M}.$$

Proof By induction on n. For the induction step, suppose we are treating the condition $\left(\inf_y \psi(x_1, \dots, x_n, y)\right) = 0$. If $\mathcal{M}$ is ω-saturated, then $\left(\inf_y \psi(a_1, \dots, a_n, y)\right) = 0$ holds in $\mathcal{M}$ if and only if there exists some $b \in M$ such that $\psi(a_1, \dots, a_n, b) = 0$ holds in $\mathcal{M}$. For the left to right direction, take $b \in M$ to satisfy the conditions $\psi(a_1, \dots, a_n, b) \leq 1/n$ for $n \geq 1$. □

7.8 Definition. Let $\mathcal{M}$ be an L-structure and let $\mathcal{N}$ be an elementary extension of $\mathcal{M}$. We call $\mathcal{N}$ an *enlargement* of $\mathcal{M}$ if it has the following property: whenever $A \subseteq M$ and $\Gamma(x_1, \dots, x_n)$ is a set of $L(A)$-conditions, if every finite subset of Γ is satisfiable in $(\mathcal{M}, a)_{a \in A}$, then the entire set Γ is satisfiable in $(\mathcal{N}, a)_{a \in A}$.

7.9 Lemma. *Every L-structure has an enlargement.*

Proof Let $\mathcal{M}$ be an L-structure and let J be a set of cardinality $\geq$ card$(L(M))$. Let I be the collection of finite subsets of J and let D be an ultrafilter on I that contains each of the sets $S_j = \{i \in I \mid j \in i\}$, for $j \in J$. Such an ultrafilter exists since this collection of sets has the finite intersection property. Let $\mathcal{N}$ be the D-ultrapower of $\mathcal{M}$, considered as an elementary extension of $\mathcal{M}$ via the diagonal map. We will show that $\mathcal{N}$ is an enlargement of $\mathcal{M}$.

Let $A \subseteq M$ and suppose $\Gamma(x_1, \dots, x_n)$ is a set of $L(A)$-conditions such that every finite subset of Γ is satisfiable in $(\mathcal{M}, a)_{a \in A}$. Let α be a function from J onto Γ. Given $i = \{j_1, \dots, j_m\} \in I$, let $(a_i^1, \dots, a_i^n)$ be any n-tuple from M that satisfies the finite subset $\{\alpha(j_1), \dots, \alpha(j_m)\}$ of Γ in $(\mathcal{M}, a)_{a \in A}$. For each $k = 1, \dots, n$ set $a_k = \big((a_i^k)_{i \in I}\big)_D$. Theorem 5.4 easily yields that $(a_1, \dots, a_n)$ satisfies Γ in $(\mathcal{N}, a)_{a \in A}$. □

7.10 Proposition. *Let $\mathcal{M}$ be an L-structure. For every infinite cardinal κ, $\mathcal{M}$ has a κ-saturated elementary extension.*

Proof By increasing κ if necessary (for example, replacing κ by κ^+) we may assume κ is regular. By induction we construct an elementary chain $(\mathcal{M}_\alpha \mid \alpha < \kappa)$ such that $\mathcal{M}_0 = \mathcal{M}$ and for each $\alpha < \kappa$, $\mathcal{M}_{\alpha+1}$ is an enlargement of $\mathcal{M}_\alpha$. (At limit ordinals we take unions.) Let $\mathcal{N}$ be the union of the chain $(\mathcal{M}_\alpha \mid \alpha < \kappa)$. By Proposition 7.2, $\mathcal{M}_\alpha \preceq \mathcal{N}$ for all $\alpha < \kappa$; in particular, $\mathcal{M} \preceq \mathcal{N}$. We claim that $\mathcal{N}$ is κ-saturated. Let A be a subset of $\mathcal{N}$ of cardinality $< \kappa$. Since κ is regular, there exists $\alpha < \kappa$ such that A is a subset of $\mathcal{M}_\alpha$. The elements of $\mathcal{N}$ needed to verify Definition 7.5 for $(\mathcal{N}, a)_{a \in A}$ can be found in $\mathcal{M}_{\alpha+1}$. □

Strongly homogeneous structures

7.11 Definition. Let $\mathcal{M}$ be an L-structure and let κ be an infinite cardinal. We say that $\mathcal{M}$ is *strongly κ-homogeneous* if the following statement holds: whenever $L(C)$ is an extension of L by constants with $\mathrm{card}(C) < \kappa$ and f, g are functions from C into M such that

$$(\mathcal{M}, f(c))_{c \in C} \equiv (\mathcal{M}, g(c))_{c \in C}$$

one has

$$(\mathcal{M}, f(c))_{c \in C} \cong (\mathcal{M}, g(c))_{c \in C}.$$

Note that an isomorphism from $(\mathcal{M}, f(c))_{c \in C}$ onto $(\mathcal{M}, g(c))_{c \in C}$ is an *automorphism* of $\mathcal{M}$ that takes $f(c)$ to $g(c)$ for each $c \in C$.

7.12 Proposition. *Let $\mathcal{M}$ be an L-structure. For every infinite cardinal κ, $\mathcal{M}$ has a κ-saturated elementary extension $\mathcal{N}$ such that each reduct of $\mathcal{N}$ to a sublanguage of L is strongly κ-homogeneous.*

Proof We may assume κ is regular, by increasing κ if necessary. Given any L-structure $\mathcal{M}$, we construct an elementary chain $(\mathcal{M}_\alpha \mid \alpha < \kappa)$ whose union has the desired properties. Let $\mathcal{M}_0 = \mathcal{M}$; for each $\alpha < \kappa$, let $\mathcal{M}_{\alpha+1}$ be an elementary extension of $\mathcal{M}_\alpha$ that is τ_α-saturated, where τ_α is a cardinal bigger than $\mathrm{card}(L)$ and bigger than the cardinality of $\mathcal{M}_\alpha$; take unions at limit ordinals. Let $\mathcal{N}$ be the union of $(\mathcal{M}_\alpha \mid \alpha < \kappa)$. By Proposition 7.2, $\mathcal{M} \preceq \mathcal{N}$. An argument such as in the proof of Proposition 7.10 shows that $\mathcal{N}$ is κ-saturated.

The fact that $\mathcal{N}$ is strongly κ-homogeneous follows from an inductive argument whose successor steps are based on the following easily proved fact:

Suppose $\mathcal{M} \preceq \mathcal{N}$ and $\mathcal{N}$ is τ-saturated, where τ is a cardinal satisfying $\tau > \mathrm{card}(L)$ and $\tau > \mathrm{card}(M)$. Let C be a set of new constants with $\mathrm{card}(C) < \tau$. Suppose f, g are functions from C into $\mathcal{M}$ such that $(\mathcal{M}, f(c))_{c \in C} \equiv (\mathcal{M}, g(c))_{c \in C}$. Then there exists an elementary embedding $T \colon \mathcal{M} \to \mathcal{N}$ such that for every $c \in C$, $T(f(c)) = g(c)$.

Finally, suppose L' is a sublanguage of L. For each $\alpha < \kappa$, the reduct of $\mathcal{M}_{\alpha+1}$ to L' is also τ_α-saturated. Hence an argument similar to the one given above for $\mathcal{N}$ shows that the reduct of $\mathcal{N}$ to L' is also strongly κ-homogeneous. □

Universal domains

7.13 Definition. Let T be a complete theory in L and let κ be an infinite cardinal number. A *κ-universal domain for T* is a κ-saturated, strongly κ-homogeneous model of T. If $\mathcal{U}$ is a κ-universal domain for T and $A \subseteq \mathcal{U}$, we will say A is *small* if $\mathrm{card}(A) < \kappa$.

By Proposition 7.12, every complete theory has a κ-universal domain for every infinite cardinal κ. Indeed, T has a model $\mathcal{U}$ such that not only $\mathcal{U}$ itself but also every reduct of $\mathcal{U}$ to a sublanguage of L is κ-saturated and strongly κ-homogeneous. Such models are needed for some arguments.

Implications

One of the subtleties of continuous first-order logic is that it is essentially a *positive* logic. In particular, there is no direct way to express an implication between conditions. This is inconvenient in applications, since many natural mathematical properties are stated using implications. However, when working in a saturated model or in all models of a theory, this obstacle can be clarified and often overcome in a natural way, which we explain here.

Let L be any signature for metric structures. For the rest of this section we fix two L-formulas $\varphi(x_1, \ldots, x_n)$ and $\psi(x_1, \ldots, x_n)$ and an L-theory T. For convenience we write x for $x_1, \ldots, x_n$.

7.14 Proposition. *Let $\mathcal{M}$ be an ω-saturated model of T. The following statements are equivalent:*
(1) For all $a \in M^n$, if $\varphi^{\mathcal{M}}(a) = 0$ then $\psi^{\mathcal{M}}(a) = 0$.
(2) $\forall \epsilon > 0\ \exists \delta > 0\ \forall a \in M^n\ (\varphi^{\mathcal{M}}(a) < \delta \ \Rightarrow\ \psi^{\mathcal{M}}(a) \leq \epsilon)$.
(3) There is an increasing, continuous function $\alpha\colon [0,1] \to [0,1]$ with $\alpha(0) = 0$ such that $\psi^{\mathcal{M}}(a) \leq \alpha(\varphi^{\mathcal{M}}(a))$ for all $a \in M^n$.

Proof (1) $\Rightarrow$ (2): Suppose (1) holds. If (2) fails, then for some $\epsilon > 0$ the set of conditions $\psi(x) \geq \epsilon$ and $\varphi(x) \leq 1/n$ for $n \geq 1$ is finitely satisfiable in $\mathcal{M}$. Since $\mathcal{M}$ realizes every finitely satisfiable set of L-conditions, there exists $a \in M^n$ such that $\psi^{\mathcal{M}}(a) \geq \epsilon$ while $\varphi^{\mathcal{M}}(a) \leq 1/n$ for all $n \geq 1$. This contradicts (1).
(2) $\Rightarrow$ (3): This follows from Proposition 2.10.
(3) $\Rightarrow$ (1): This is trivial. □

The next results are variants of the previous one, with essentially the same proof. In stating them we translate (2) and (3) using conditions of continuous logic, as follows.

The statement in (2) holds for a given L-structure $\mathcal{M}$ and a given ϵ, δ if and only if the L-condition

$$\sup_x \min\left(\delta \dot{-} \varphi(x), \psi(x) \dot{-} \epsilon\right) = 0$$

is true in $\mathcal{M}$.

The statement in (3) holds for a given $\mathcal{M}$ and a given α if and only if the L-condition

$$\sup_x \left(\psi(x) \dot{-} \alpha(\varphi(x))\right) = 0$$

is true in $\mathcal{M}$.

7.15 Proposition. *The following statements are equivalent.*
(1) For all $\mathcal{M} \models T$ *and all* $a \in M^n$*, if* $\varphi^{\mathcal{M}}(a) = 0$ *then* $\psi^{\mathcal{M}}(a) = 0$.
(2) For all $\epsilon > 0$ *there exists* $\delta > 0$ *such that*

$$T \models \sup_x \min\big(\delta \dot{-} \varphi(x), \psi(x) \dot{-} \epsilon\big) = 0.$$

(3) There is an increasing, continuous function $\alpha\colon [0,1] \to [0,1]$ *with* $\alpha(0) = 0$ *such that*

$$T \models \sup_x \big(\psi(x) \dot{-} \alpha(\varphi(x))\big) = 0.$$

Proof It is obvious that (2) and (3) imply (1). If (2) fails, then for some $\epsilon > 0$ the set of conditions $\psi(x) \geq \epsilon$ and $\varphi(x) \leq 1/n$ for $n \geq 1$ is finitely satisfiable in models of T. The Compactness Theorem yields $\mathcal{M} \models T$ and $a \in M^n$ such that $\psi^{\mathcal{M}}(a) \geq \epsilon$ while $\varphi^{\mathcal{M}}(a) \leq 1/n$ for all $n \geq 1$. This contradicts (1).

Now assume that (1) and (equivalently) (2) hold. Define a function Δ on $(0,1]$ by setting each $\Delta(\epsilon)$ to be half the supremum of all $\delta \in (0,1]$ for which $T \models \sup_x \min\big(\delta \dot{-} \varphi(x), \psi(x) \dot{-} \epsilon\big) = 0$. Evidently Δ is an increasing function and (since (2) holds) it takes its values in $(0,1]$. For each $\mathcal{M} \models T$ and each $\epsilon \in (0,1]$ we have

$$\forall a \in M^n \left(\varphi^{\mathcal{M}}(a) \leq \Delta(\epsilon) \Rightarrow \psi^{\mathcal{M}}(a) \leq \epsilon\right).$$

Now we use the argument in Remark 2.12 to obtain an increasing, continuous function $\alpha\colon [0,1] \to [0,1]$ with $\alpha(0) = 0$ such that for each $\mathcal{M} \models T$ we have

$$\forall a \in M^n \left(\psi^{\mathcal{M}}(a) \leq \alpha(\varphi^{\mathcal{M}}(a))\right).$$

This proves that statement (3) holds for this α. □

7.16 Corollary. *Let* $\mathcal{C}$ *be the set of all models* $\mathcal{M}$ *of* T *that satisfy the requirement*

$$\forall a \in M^n \ \left[\varphi^{\mathcal{M}}(a) = 0 \ \Rightarrow \ \psi^{\mathcal{M}}(a) = 0\right].$$

Then the following statements are equivalent:
(1) $\mathcal{C}$ *is axiomatizable.*
(2) For each $\epsilon > 0$ *there exists* $\delta > 0$ *such that the* L*-condition*

$$\sup_x \min\big(\delta \dot{-} \varphi(x), \psi(x) \dot{-} \epsilon\big) = 0$$

is true in all members of $\mathcal{C}$.

(3) There exists an increasing, continuous function $\alpha\colon [0,1] \to [0,1]$ *with* $\alpha(0) = 0$ *such that the L-condition*

$$\sup_x \bigl(\psi(x) \dot{-} \alpha(\varphi(x))\bigr) = 0$$

is true in all members of $\mathcal{C}$.

Proof If (1) holds, apply the previous result to a theory T' that axiomatizes $\mathcal{C}$. If (2) or (3) holds, $\mathcal{C}$ can be axiomatized by adding conditions of the form displayed to T. □

8 Spaces of types

In this section we consider a fixed signature L for metric structures and a fixed L-theory T. Until further notice in this section we assume that T is a complete theory.

Suppose that $\mathcal{M}$ is a model of T and $A \subseteq M$. Denote the $L(A)$-structure $(\mathcal{M}, a)_{a \in A}$ by $\mathcal{M}_A$, and set T_A to be the $L(A)$-theory of $\mathcal{M}_A$. Note that any model of T_A is isomorphic to a structure of the form $(\mathcal{N}, a)_{a \in A}$, where $\mathcal{N}$ is a model of T.

8.1 Definition. Let T_A be as above and let $x_1, \ldots, x_n$ be distinct variables.

A set p of $L(A)$-conditions with all free variables among $x_1, \ldots, x_n$ is called an *n-type over A* if there exists a model $(\mathcal{M}, a)_{a \in A}$ of T_A and elements $e_1, \ldots, e_n$ of M such that p is the set of all $L(A)$-conditions $E(x_1, \ldots, x_n)$ for which $\mathcal{M}_A \models E[e_1, \ldots, e_n]$.

When this relationship holds, we denote p by $\mathrm{tp}_{\mathcal{M}}(e_1, \ldots, e_n/A)$ and we say that $(e_1, \ldots, e_n)$ *realizes p in $\mathcal{M}$*. (The subscript $\mathcal{M}$ will be omitted if doing so causes no confusion; A will be omitted if it is empty.)

The collection of all such n-types over A is denoted by $S_n(T_A)$, or simply by $S_n(A)$ if the context makes the theory T_A clear.

8.2 Remark. Let $\mathcal{M}, A$ be as above, and let e, e' be n-tuples from M.
(1) $\mathrm{tp}_{\mathcal{M}}(e/A) = \mathrm{tp}_{\mathcal{M}}(e'/A)$ if and only if $(\mathcal{M}_A, e) \equiv (\mathcal{M}_A, e')$.
(2) If $\mathcal{M} \preceq \mathcal{N}$, then $\mathrm{tp}_{\mathcal{M}}(e/A) = \mathrm{tp}_{\mathcal{N}}(e/A)$.

8.3 Remark. Suppose $\mathcal{M}$ is a κ-saturated L-structure. It is easy to show that for any subset $A \subseteq M$ with $\mathrm{card}(A) < \kappa$, every type in $S_n(T_A)$ is realized in $\mathcal{M}$. Indeed, this property (even just with $n = 1$) is equivalent to κ-saturation of $\mathcal{M}$.

The logic topology on types

Fix T_A as above. If $\varphi(x_1,\dots,x_n)$ is an $L(A)$-formula and $\epsilon > 0$, we let $[\varphi < \epsilon]$ denote the set

$$\{q \in S_n(T_A) \mid \text{for some } 0 \leq \delta < \epsilon \text{ the condition } (\varphi \leq \delta) \text{ is in } q\ \}.$$

8.4 Definition. The *logic topology* on $S_n(T_A)$ is defined as follows. If p is in $S_n(T_A)$, the basic open neighborhoods of p are the sets of the form $[\varphi < \epsilon]$ for which the condition $\varphi = 0$ is in p and $\epsilon > 0$.

Note that the logic topology is Hausdorff. Indeed, if p, q are distinct elements of $S_n(T_A)$, then there exists an $L(A)$-formula $\varphi(x_1,\dots,x_n)$ such that the condition $\varphi = 0$ is in one of the types but not the other. Therefore, for some positive r the condition $\varphi = r$ is in that other type. Taking $\epsilon = r/2 > 0$ we see that $[\varphi < \epsilon]$ and $[(r \dot- \varphi) < \epsilon]$ are disjoint basic open sets for the logic topology, one containing p and the other containing q.

It is also useful to introduce notation such as the following (where $\varphi(x_1,\dots,x_n)$ is an $L(A)$-formula and $\epsilon \geq 0$):

$$[\varphi \leq \epsilon] = \{q \in S_n(T_A) \mid \text{the condition } (\varphi \leq \epsilon) \text{ is in } q\ \}.$$

Each set $[\varphi \leq \epsilon]$ is closed in the logic topology; indeed, its complement is $\emptyset$ if $\epsilon \geq 1$ and it is $[1 \dot- \varphi < \delta]$ if $\epsilon < 1$ and $\delta = 1 - \epsilon$.

8.5 Lemma. *The closed subsets of $S_n(T_A)$ for the logic topology are exactly the sets of the form $C_\Gamma = \{p \in S_n(T_A) \mid \Gamma(x_1,\dots,x_n) \subseteq p\}$ where $\Gamma(x_1,\dots,x_n)$ is a set of $L(A)$-conditions.*

Proof Given such a set $\Gamma(x_1,\dots,x_n)$, note that C_Γ is the intersection of all sets $[\varphi \leq 0]$ where $\varphi = 0$ is any condition in Γ. Hence C_Γ is closed. Conversely, suppose C is a subset of $S_n(T_A)$ that is closed in the logic topology and let p be any element of $S_n(T_A) \setminus C$. By the definition of the logic topology there exists an $L(A)$-condition $\varphi = 0$ in p and $\epsilon > 0$ such that $[\varphi < \epsilon]$ is disjoint from C. Without loss of generality we may assume $\epsilon \leq 1$. Then the closed set $[(\epsilon \dot- \varphi) \leq 0]$ contains C and does not have p as an element. To represent C in the desired form it suffices to take Γ to be the set of all conditions of the form $(\epsilon \dot- \varphi) = 0$ that arise in this way. □

8.6 Proposition. *For any $n \geq 1$, $S_n(T_A)$ is compact with respect to the logic topology.*

Proof In light of the preceding discussion, one sees that this is just a restatement of the Compactness Theorem (Corollary 5.12). □

The d-metric on types

Let T_A be as above. For each $n \geq 1$ we define a natural metric on $S_n(T_A)$; it is induced as a quotient of the given metric d on M^n, where $(\mathcal{M}, a)_{a\in A}$ is a suitable model of T_A, so we also denote this metric on types by d.

To define this metric, let $\mathcal{M}_A = (\mathcal{M}, a)_{a\in A}$ be any model of T_A in which each type in $S_n(T_A)$ is realized, for each $n \geq 1$. (Such a model exists by Proposition 7.10.) Let (M, d) be the underlying metric space of $\mathcal{M}$. For $p, q \in S_n(T_A)$ we define $d(p, q)$ to be

$$\inf\Big\{ \max_{1\leq j\leq n} d(b_j, c_j) \;\Big|\; \mathcal{M}_A \models p[b_1, \ldots, b_n],\;\; \mathcal{M}_A \models q[c_1, \ldots, c_n] \Big\}.$$

Note that this expression for $d(p, q)$ does not depend on $\mathcal{M}_A$, since $\mathcal{M}_A$ realizes every type of a $2n$-tuple $(b_1, \ldots, b_n, c_1, \ldots, c_n)$ over A. It follows easily that d is a pseudometric on $S_n(T_A)$. Note that if $p, q \in S_n(A)$, then by the Compactness Theorem and our assumptions about $\mathcal{M}_A$, there exist realizations $(b_1, \ldots, b_n)$ of p and $(c_1, \ldots, c_n)$ of q in $\mathcal{M}_A$, such that $\max_j d(b_j, c_j) = d(p, q)$. In particular, if $d(p, q) = 0$, then $p = q$; so d is indeed a metric on $S_n(T_A)$.

8.7 Proposition. *The d-topology is finer than the logic topology on $S_n(T_A)$.*

Proof This follows from the uniform continuity of formulas. □

8.8 Proposition. *The metric space $(S_n(T_A), d)$ is complete.*

Proof Let $(p_k)_{k\geq 1}$ be a Cauchy sequence in $(S_n(T_A), d)$. Without loss of generality we may assume $d(p_k, p_{k+1}) \leq 2^{-k}$ for all k; that is, for completeness it suffices to show that every such Cauchy sequence has a limit. Let $\mathcal{N}$ be an ω-saturated and strongly ω-homogeneous model of T_A. Without loss of generality we may assume that $\mathcal{N} = \mathcal{M}_A$ for some model $\mathcal{M}$ of T. Using our saturation and homogeneity assumptions, we see that for any $a \in M^n$ realizing p_k there exists $b \in M^n$ realizing p_{k+1} such that $d(a, b) = d(p_k, p_{k+1})$. Therefore, proceeding inductively we may generate a sequence (b_k) in M^n such that b_k realizes p_k and $d(b_k, b_{k+1}) = d(p_k, p_{k+1}) \leq 2^{-k}$ for all k. This implies that (b_k) is a Cauchy sequence in M^n so it converges in M^n to some $b \in M^n$. It follows that the type realized by b in $(\mathcal{M}, a)_{a\in A}$ is the limit of the sequence (p_k) in the metric space $(S_n(T_A), d)$. □

Functions on type spaces defined by formulas

Let T_A be as above. Let $\mathcal{M}_A = (\mathcal{M}, a)_{a\in A}$ be any model of T_A in which each type in $S_n(T_A)$ is realized, for each $n \geq 1$.

Let $\varphi(x_1, \dots, x_n)$ be any $L(A)$-formula. For each type $p \in S_n(T_A)$ we let $\widetilde{\varphi}(p)$ denote the unique real number $r \in [0, 1]$ for which the condition $\varphi = r$ is in p. Equivalently, $\widetilde{\varphi}(p) = \varphi^{\mathcal{M}}(b)$ when b is any realization of p in $\mathcal{M}_A$.

8.9 Lemma. *Let $\varphi(x_1, \dots, x_n)$ be any $L(A)$-formula. The function $\widetilde{\varphi}\colon S_n(T_A) \to [0, 1]$ is continuous for the logic topology and uniformly continuous for the d-metric distance on $S_n(T_A)$.*

Proof For any $r \in [0, 1]$ and $\epsilon > 0$, note that

$$\widetilde{\varphi}^{-1}(r - \epsilon, r + \epsilon) = [|\varphi \dot{-} r| < \epsilon].$$

This shows that $\widetilde{\varphi}$ is continuous for the logic topology. For the uniform continuity, use Theorem 3.5 to obtain a modulus of uniform continuity Δ_φ for $\varphi^{\mathcal{M}}$ on M^n. It is easy to show that Δ_φ is a modulus of uniform continuity for $\widetilde{\varphi}$. Indeed, suppose $\epsilon \in (0, 1]$ and let $\delta = \Delta_\varphi(\epsilon)$. Suppose $p, q \in S_n(T_A)$ have $d(p, q) < \delta$. Suppose $r = \widetilde{\varphi}(p)$ and $s = \widetilde{\varphi}(q)$. We need to show $|r - s| \leq \epsilon$. Choose $a, b \in M^n$ to realize p, q respectively in $\mathcal{M}_A$ with $d(a, b) = d(p, q)$. Our choice of Δ_φ ensures that $|r - s| = |\varphi^{\mathcal{M}}(a) - \varphi^{\mathcal{M}}(b)| \leq \epsilon$, as desired. □

8.10 Proposition. *For any function $\Phi\colon S_n(T_A) \to [0, 1]$ the following statements are equivalent:*
(1) Φ is continuous for the logic topology on $S_n(T_A)$;
(2) There is a sequence $(\varphi_k(x_1, \dots, x_n))_{k\geq 1}$ of $L(A)$-formulas such that $(\widetilde{\varphi_k})_{k\geq 1}$ converges to Φ uniformly on $S_n(T_A)$;
(3) Φ is continuous for the logic topology and uniformly continuous for the d-metric on $S_n(T_A)$.

Proof (1) ⇒ (2): Note that the set of functions of the form $\widetilde{\varphi}$, where $\varphi(x_1, \dots, x_n)$ is an $L(A)$-formula, separates points of $S_n(T_A)$ and is closed under the pointwise lattice operations max and min. Applying the lattice version of the Stone-Weierstrass Theorem to the compact, Hausdorff space $S_n(T_A)$ with the logic topology yields (2).

(2) ⇒ (3): Each $\widetilde{\varphi}$ is continuous for the logic topology and uniformly continuous for the d-metric on $S_n(T_A)$. These properties are preserved under uniform convergence.

(3) ⇒ (1): trivial. □

There are many connections between the type spaces $S_n(T_A)$ as n and A vary. We conclude this subsection with a result that summarizes some of them.

8.11 Proposition. *Let $\mathcal{M} \models T$ and $A \subseteq B \subseteq M$. Let π be the restriction map from $S_n(T_B)$ to $S_n(T_A)$(defined by letting $\pi(p)$ be the set of $L(A)$-conditions in p). Then:*
(1) π is surjective;
(2) π is continuous (hence closed) for the logic topologies;
(3) π is contractive (hence uniformly continuous) for the d-metrics;
(4) if A is d-dense in B, then π is a homeomorphism for the logic topologies and a surjective isometry for the d-metrics.

Proof (1) Let $p \in S_n(T_A)$. The set p^+ of $L(A)$-formulas is finitely satisfied in $(\mathcal{M}, b)_{b\in B}$, so p itself is satisfied in some elementary extension of $(\mathcal{M}, b)_{b\in B}$, by the Compactness Theorem. If $(e_1, \ldots, e_n)$ realizes p in such an elementary extension, then $p = \pi(\text{tp}(e_1, \ldots, e_n/B)$.
(2) If $\varphi(x_1, \ldots, x_n)$ is any $L(A)$-formula and $\epsilon > 0$, then π obviously maps $[\varphi < \epsilon]$ as a basic neighborhood in $S_n(T_B)$ into $[\varphi < \epsilon]$ as a basic neighborhood in $S_n(T_A)$. Therefore π is continuous for the logic topologies; hence π is also a closed map, since those topologies are compact and Hausdorff.
(3) Any realization of $p \in S_n(T_B)$ is also a realization of $\pi(p)$.
(4) This follows from the fact that formulas define uniformly continuous functions. Hence any $L(B)$-formula can be uniformly approximated by a sequence of $L(A)$-formulas, when A is dense in B. In particular, if $p \in S_n(T_B)$, then any realization of $\pi(p)$ is a realization of p itself. □

Types over an arbitrary theory

For the rest of this section we consider a theory T that is satisfiable but not necessarily complete. For some purposes one needs to consider the spaces of types over $\emptyset$ that are consistent with T.

Fix $n \geq 0$ and let $x_1, \dots, x_n$ be distinct variables. We let $S_n(T)$ denote the set of all n-types over $\emptyset$ of the form $\mathrm{tp}_{\mathcal{M}}(e_1, \dots, e_n)$ where $\mathcal{M}$ is any model of T and $e_1, \dots, e_n$ are in M. We equip $S_n(T)$ with the *logic topology*, whose basic open sets are of the form

$$[\varphi < \epsilon] = \{q \in S_n(T) \ \mid \text{condition } (\varphi \leq \delta) \text{ is in } q \text{ for some } 0 \leq \delta < \epsilon\}$$

where $\epsilon > 0$ and $\varphi(x_1, \dots, x_n)$ is an L-formula. (Note that $S_0(T)$ is simply the space of all complete L-theories that extend T.)

For each L-formula $\varphi(x_1, \dots, x_n)$ we define $\widetilde{\varphi}\colon S_n(T) \to [0,1]$ by setting $\widetilde{\varphi}(p)$ to be the unique real number $r \in [0,1]$ for which the condition $\varphi = r$ is in p, for each $p \in S_n(T)$. Equivalently, $\widetilde{\varphi}(p) = \varphi^{\mathcal{M}}(a)$ where a is any realization of p in any $\mathcal{M} \models T$. The proof of Lemma 8.9 shows that each $\widetilde{\varphi}$ is continuous as a function into $[0,1]$ from $S_n(T)$ with the logic topology.

8.12 Theorem. *Let T be any satisfiable L-theory; $S_n(T)$ equipped with the logic topology has the following properties:*

(1) *The closed subsets of $S_n(T)$ are exactly the sets of the form*

$$C_\Gamma = \{p \in S_n(T) \mid \Gamma(x_1, \dots, x_n) \subseteq p\}$$

where $\Gamma(x_1, \dots, x_n)$ is a set of L-conditions.

(2) *$S_n(T)$ is a compact, Hausdorff space.*

(3) *$\Phi\colon S_n(T) \to [0,1]$ is continuous if and only if there is a sequence $(\varphi_k(x_1, \dots, x_n))_{k \geq 1}$ of L-formulas such that $(\widetilde{\varphi_k})_{k \geq 1}$ converges to Φ uniformly on $S_n(T)$.*

Proof The proofs of Lemma 8.5 and Propositions 8.6 and 8.10 apply without change to the current situation. □

9 Definability in metric structures

In this section we discuss some issues around *definability*, which is arguably the central topic in model theory and its applications.

Let $\mathcal{M}$ be a metric structure, and L a signature for $\mathcal{M}$. First we consider definability of predicates $P\colon M^n \to [0,1]$. Then we use definable predicates to discuss definability of subsets of M^n and functions from M^n into M.

Let A be any subset of M, which we think of as a set of possible parameters to use in definitions.

Definable predicates

9.1 Definition. A predicate $P\colon M^n \to [0,1]$ is definable in $\mathcal{M}$ over A if and only if there is a sequence $(\varphi_k(x) \mid k \geq 1)$ of $L(A)$-formulas such that the predicates $\varphi_k^{\mathcal{M}}(x)$ converge to $P(x)$ uniformly on M^n; *i.e.*,

$$\forall \epsilon > 0\ \exists N\ \forall k \geq N\ \forall x \in M^n\ \left(|\varphi_k^{\mathcal{M}}(x) - P(x)| \leq \epsilon\right).$$

In other words, a predicate is definable over A if it is in the uniform closure of the set of functions from M^n to $[0,1]$ that are obtained by interpreting $L(A)$-formulas in $\mathcal{M}$. We will give various results to show why we think this is the "right" notion of definability for predicates in metric structures.

9.2 Remark. Suppose $P\colon M^n \to [0,1]$ is definable in $\mathcal{M}$ over A. By taking $A_0 \subseteq A$ to be the set of elements of A whose names appear in the sequence of $L(A)$-formulas $(\varphi_k(x) \mid k \geq 1)$ witnessing that P is definable, we get a countable set A_0 such that P is definable in $\mathcal{M}$ over A_0. In contrast to what happens in ordinary first-order logic, it need not be possible to do this with a *finite* set of parameters.

It can be useful to have a specific representation of definable predicates. To do this we broaden our perspective of "connectives". Consider the product space $[0,1]^{\mathbb{N}}$ of infinite sequences; this is a compact metrizable space; for example, its topology is given by the metric ρ defined by

$$\rho\left((a_k),(b_k)\right) = \sum_{k=0}^{\infty} 2^{-k}|a_k - b_k|$$

for any pair of sequences $(a_k \mid k \in \mathbb{N})$ and $(b_k \mid k \in \mathbb{N})$ in $[0,1]^{\mathbb{N}}$. In representing definable predicates we will regard any continuous function $u\colon [0,1]^{\mathbb{N}} \to [0,1]$ as a kind of connective.

9.3 Proposition. *Let $\mathcal{M}$ be an L-structure with $A \subseteq M$, and suppose $P\colon M^n \to [0,1]$ is a predicate. Then P is definable in $\mathcal{M}$ over A if and only if there is a continuous function $u\colon [0,1]^{\mathbb{N}} \to [0,1]$ and $L(A)$-formulas $(\varphi_k \mid k \in \mathbb{N})$ such that for all $x \in M^n$*

$$P(x) = u\left(\varphi_k^{\mathcal{M}}(x) \mid k \in \mathbb{N}\right).$$

Proof Suppose P has the specified form. Fix $\epsilon > 0$; we need to show that P can be approximated uniformly to within ϵ by the interpretation of some $L(A)$-formula.

Since $([0,1]^{\mathbb{N}}, \rho)$ is compact, the function u is uniformly continuous with respect to ρ. Hence there exists $m \in \mathbb{N}$ such that

$$|u((a_k)) - u((b_k))| \leq \epsilon$$

holds whenever $(a_k \mid k \in \mathbb{N})$ and $(b_k \mid k \in \mathbb{N})$ are sequences such that $a_k = b_k$ for all $k = 0, \dots, m$. Let $u_m \colon [0,1]^{m+1} \to [0,1]$ be the continuous function obtained from u by setting

$$u_m(a_0, \dots, a_m) = u(a_0, \dots, a_m, 0, 0, 0, \dots)$$

for all $a_0, \dots, a_m \in \mathbb{N}$. Let $\varphi(x)$ be the $L(A)$-formula given by

$$\varphi(x) = u_m(\varphi_0(x), \dots, \varphi_m(x)).$$

Then we have immediately that

$$|P(x) - \varphi^{\mathcal{M}}(x)| = |u\left(\varphi_k^{\mathcal{M}}(x) \mid k \in \mathbb{N}\right) - u_m\left(\varphi_0^{\mathcal{M}}(x), \dots, \varphi_m^{\mathcal{M}}(x)\right)| \leq \epsilon$$

for all $x \in M^n$. Hence P is definable over A in $\mathcal{M}$.

For the converse direction, assume that $P \colon M^n \to [0,1]$ is definable over A in $\mathcal{M}$. For each $k \in \mathbb{N}$ let $\varphi_k(x)$ be an $L(A)$-formula such that

$$|\varphi_k^{\mathcal{M}}(x) - P(x)| \leq 2^{-k}$$

holds for all $x \in M^n$.

Now consider the set C of all sequences $(a_k \mid k \in \mathbb{N})$ in $[0,1]^{\mathbb{N}}$ such that $|a_k - a_l| \leq 2^{-N}$ holds whenever $N \in \mathbb{N}$ and $k, l \geq N+1$. Each sequence (a_k) in C is a Cauchy sequence in $[0,1]$, so it converges to a limit that we denote by $\lim(a_k)$. It is easy to check that C is a closed subset of $[0,1]^{\mathbb{N}}$ and that the sequence $(\varphi_k^{\mathcal{M}}(x) \mid k \in \mathbb{N})$ is in C for every $x \in M^n$. Moreover, it is easy to check that the function $\lim \colon C \to [0,1]$ is continuous with respect to the restriction of the product topology on $[0,1]^{\mathbb{N}}$ to C. By the Tietze Extension Theorem, there is a continuous function $u \colon [0,1]^{\mathbb{N}} \to [0,1]$ that agrees with lim on C. From this we conclude immediately that

$$P(x) = u\left(\varphi_k^{\mathcal{M}}(x) \mid k \in \mathbb{N}\right)$$

for all $x \in M^n$, as desired. □

9.4 Remark. Note that in our proof of Proposition 9.3, the continuous function $u \colon [0,1]^{\mathbb{N}} \to [0,1]$, in terms of which the definable predicate P was represented, is completely independent of P. It is not too difficult to give a constructive proof of this result, in which u is described concretely, and doing so can be useful. See, for example, the *forced limit function* discussed in [6, section 3].

9.5 Remark. Proposition 9.3 shows one way to represent definable predicates so that they become meaningful in every L-structure. This suggests how the notion of L-formula could be expanded by allowing continuous *infinite* connectives, without expanding the notion of *definability* for predicates, in order to have an exact correspondence between *formulas* and *definable predicates*. There is the complication that, as noted above, a definable predicate may depend on *infinitely many* of the parameters that are used in its definition. We will not explore this direction here.

9.6 Lemma. *Suppose $P\colon M^n \to [0,1]$ is definable in $\mathcal{M}$ over A and consider $\mathcal{N} \preceq \mathcal{M}$ with $A \subseteq N$. Then $\inf_x P(x)$ and $\sup_x P(x)$ have the same value in $\mathcal{N}$ as in $\mathcal{M}$.*

Proof Let $(\varphi_k \mid k \geq 1)$ be $L(A)$-formulas such that for all $x \in M^n$ and all $k \geq 1$ we have

$$|P(x) - \varphi_k^{\mathcal{M}}(x)| \leq \frac{1}{k}.$$

Since $\mathcal{N} \preceq \mathcal{M}$, we also have for all $x \in N^n$ and all $k \geq 1$

$$|P(x) - \varphi_k^{\mathcal{N}}(x)| \leq \frac{1}{k}.$$

We conclude that

$$\begin{aligned} \inf_x P(x) &= \lim_{k\to\infty} \inf_x \varphi_k^{\mathcal{M}}(x) && (\text{inf over } x \in M^n) \\ &= \lim_{k\to\infty} \inf_x \varphi_k^{\mathcal{N}}(x) && (\text{inf over } x \in N^n) \\ &= \inf_x P(x) && (\text{inf over } x \in N^n) \end{aligned}$$

as claimed. □

This Lemma is a special case of a more general result:

9.7 Proposition. *Let $P_i\colon M^n \to [0,1]$ be definable in $\mathcal{M}$ over A for $i = 1,\dots,m$ and consider $\mathcal{N} \preceq \mathcal{M}$ with $A \subseteq N$. Let Q_i be the restriction of P_i to N^n for each i. Then $(\mathcal{N}, Q_1,\dots,Q_m) \preceq (\mathcal{M}, P_1,\dots,P_m)$.*

Proof This is proved using an elaboration of the ideas above. The proof is by induction on formulas, using the tools concerning uniform convergence that were developed in the appendix to section 2. □

9.8 Proposition. *Let $P\colon M^n \to [0,1]$ be definable in $\mathcal{M}$ over A and consider an elementary extension $\mathcal{N}$ of $\mathcal{M}$. There is a unique predicate $Q\colon N^n \to [0,1]$ such that Q is definable in $\mathcal{N}$ over A and P is the restriction of Q to M^n. This predicate satisfies $(\mathcal{M}, P) \preceq (\mathcal{N}, Q)$.*

Proof Let $(\varphi_k(x) \mid k \geq 1)$ be a sequence of $L(A)$-formulas such that the functions $(\varphi_k^{\mathcal{M}} \mid k \geq 1)$ converge uniformly to P on M^n. Note that for any $k, l \geq 1$ we have

$$\sup\{|\varphi_k^{\mathcal{N}}(b) - \varphi_l^{\mathcal{N}}(b)| \mid b \in N^n\} = \sup\{|\varphi_k^{\mathcal{M}}(a) - \varphi_l^{\mathcal{M}}(a)| \mid a \in M^n\}.$$

Therefore the functions $(\varphi_k^{\mathcal{N}} \mid k \geq 1)$ converge uniformly on N^n to some function $Q\colon N^n \to [0,1]$. Evidently Q extends P and the construction of Q ensures that it is definable in $\mathcal{N}$ over A. The last statement follows from Proposition 9.7.

It remains to prove Q is unique. Suppose Q_1, Q_2 are predicates definable in $\mathcal{M}$ over A whose restriction to M^n equals P. Applying Proposition 9.7 we conclude that $(\mathcal{N}, Q_1, Q_2)$ is an elementary extension of $(\mathcal{M}, P, P)$ and hence $\sup\{|Q_1(x) - Q_2(x)| \mid x \in N^n\} = \sup\{|P(x) - P(x)| \mid x \in M^n\} = 0$. Therefore $Q_1 = Q_2$. □

The following result gives a useful and conceptually appealing characterization of definable predicates. As above we have an L-structure $\mathcal{M}$ with $A \subseteq M$ and we set $\mathcal{M}_A = (\mathcal{M}, a)_{a \in M}$; let $T = \mathrm{Th}(\mathcal{M})$ be the complete L-theory of which $\mathcal{M}$ is a model and let $T_A = \mathrm{Th}(\mathcal{M}_A)$.

9.9 Theorem. *Let $P\colon M^n \to [0,1]$ be a function. Then P is a predicate definable in $\mathcal{M}$ over A if and only if there exists $\Phi\colon S_n(T_A) \to [0,1]$ that is continuous with respect to the logic topology on $S_n(T_A)$ such that $P(a) = \Phi(\mathrm{tp}_{\mathcal{M}}(a/A))$ for all $a \in M^n$.*

Proof First suppose that there is a continuous $\Phi\colon S_n(T_A) \to [0,1]$ such that $P(a) = \Phi(\mathrm{tp}_{\mathcal{M}}(a/A))$ for all $a \in M^n$. By Proposition 8.10 there is a sequence $(\varphi_k(x) \mid k \geq 1)$ of $L(A)$-formulas such that the functions $(\widetilde{\varphi}_k \mid k \geq 1)$ converge uniformly to Φ on $S_n(T_A)$. For any $a \in M^n$ let $p = \mathrm{tp}_{\mathcal{M}}(a/A)$ and note that $|\varphi_k^{\mathcal{M}}(a) - P(a)| = |\widetilde{\varphi}_k(p) - \Phi(p)|$. Therefore the functions $(\varphi_k^{\mathcal{M}} \mid k \geq 1)$ converge uniformly to P on M^n, from which it follows that P is a predicate (*i.e.*, that it is uniformly continuous) and that it is definable in $\mathcal{M}$ over A.

For the converse, suppose $(\varphi_k(x) \mid k \geq 1)$ is a sequence of $L(A)$-formulas such that the functions $(\varphi_k^{\mathcal{M}} \mid k \geq 1)$ converge uniformly to

P on M^n. Let $\mathcal{N}$ be a κ-saturated elementary extension of $\mathcal{M}$, where $\kappa > \operatorname{card}(A)$. Arguing as in the proof of Proposition 9.8 we see that the functions $(\varphi_k^{\mathcal{N}} \mid k \geq 1)$ converge uniformly on N^n to some function $Q\colon N^n \to [0,1]$. Evidently Q extends P and the construction of Q ensures that it is definable in $\mathcal{N}$ over A. Let p be any type in $S_n(T_A)$. Define $\Phi(p) = Q(b)$ where $b \in N^n$ realizes p; since $Q(b)$ is the limit of $(\varphi_k^{\mathcal{N}}(b) \mid k \geq 1)$ and $(\varphi_k \mid k \geq 1)$ are $L(A)$-formulas, the value of $Q(b)$ depends only on the type of b over A. Moreover, our construction of Q ensures that Φ is the uniform limit of the functions $(\widetilde{\varphi}_k \mid k \geq 1)$ on $S_n(T_A)$. Therefore Φ is continuous with respect to the logic topology by Proposition 8.10. For any $a \in M^n$ we have $P(a) = Q(a) = \Phi(\operatorname{tp}_{\mathcal{M}}(a/A))$ as desired. □

The next result provides a characterization of definability for predicates in saturated models that proves to be technically helpful in many situations. If $\mathcal{M}$ is an L-structure and $A \subseteq M$, a subset $S \subseteq M^n$ is called *type-definable* in $\mathcal{M}$ over A if there is a set $\Sigma(x_1, \ldots, x_n)$ of $L(A)$-formulas such that for any $a \in M^n$ we have $a \in S$ if and only if $\varphi^{\mathcal{M}}(a) = 0$ for every $\varphi \in \Sigma$. In this case we will say S is *type-defined by* Σ.

9.10 Corollary. *Let $\mathcal{M}$ be a κ-saturated L-structure and $A \subseteq M$ with* $\operatorname{card}(A) < \kappa$*; let $P\colon M^n \to [0,1]$ be a function. Then P is a predicate definable in $\mathcal{M}$ over A if and only if the sets $\{a \in M^n \mid P(a) \leq r\}$ and $\{a \in M^n \mid P(a) \geq r\}$ are type-definable in $\mathcal{M}$ over A for every $r \in [0,1]$.*

Proof Suppose P is a predicate definable in $\mathcal{M}$ over A. Using Theorem 9.9 we get a continuous function $\Phi\colon S_n(T_A) \to [0,1]$ such that $P(a) = \Phi(\operatorname{tp}_{\mathcal{M}}(a/A))$ for all $a \in M^n$. For $r \in [0,1]$, $\Phi^{-1}([0,r])$ is a closed subset of $S_n(A)$ for the logic topology. By Lemma 8.5, it is of the form $\{p \in S_n(T_A) \mid \Gamma(x_1, \ldots, x_n) \subseteq p\}$ where $\Gamma(x_1, \ldots, x_n)$ is some set of $L(A)$-conditions. It follows that $\{a \in M^n \mid P(a) \leq r\}$ is type-defined by Γ. A similar argument applies to $\{a \in M^n \mid P(a) \geq r\}$.

Conversely, suppose P is a function such that $\{a \in M^n \mid P(a) \leq r\}$ and $\{a \in M^n \mid P(a) \geq r\}$ are type-definable in $\mathcal{M}$ over A for every $r \in [0,1]$. This allows us to define $\Phi\colon S_n(T_A) \to [0,1]$ by setting $\Phi(p) = P(a)$ whenever $p \in S_n(T_A)$ and $a \in M^n$ realizes p in $\mathcal{M}_A$. Note that such an a exists for every p because of our saturation assumption. Further, our type-definability assumption ensures that Φ is well defined and, moreover, that it is continuous for the logic topology. It follows by Theorem 9.9 that P is in $\mathcal{M}$ over A. □

As another corollary to Theorem 9.9 we get a characterization of definability for predicates that is a generalization of the Theorem of Svenonius from ordinary model theory. Note that we here assume that the given function P is a predicate, *i.e.*, that it is uniformly continuous; this allows us to consider the expansion $(\mathcal{M}, P)$ as a metric structure.

9.11 Corollary. *Let $\mathcal{M}$ be an L-structure with $A \subseteq M$, and suppose $P\colon M^n \to [0,1]$ is a predicate. Then P is definable in $\mathcal{M}$ over A if and only if whenever $(\mathcal{N}, Q) \succeq (\mathcal{M}, P)$, the predicate Q is invariant under all automorphisms of $\mathcal{N}$ that leave A fixed pointwise.*

Proof For the left to right direction, assume that P is definable in $\mathcal{M}$ over A. So there is a sequence of $L(A)$-formulas $(\varphi_k \mid \kappa \geq 1)$ such that P is the uniform limit of $(\varphi_k^{\mathcal{M}} \mid k \geq 1)$ on M^n. Suppose $(\mathcal{N}, Q)$ is any elementary extension of $(\mathcal{M}, P)$. As discussed in the previous proof, we have that Q is the uniform limit of $(\varphi_k^{\mathcal{N}} \mid k \geq 1)$ N^n. Since each function $\varphi_k^{\mathcal{N}}$ is the interpretation of an $L(A)$-formula, it must be invariant under all automorphisms of $\mathcal{N}$ that leave A pointwise fixed. Hence the same is true of its uniform limit Q.

For the right to left direction, let $(\mathcal{N}, Q) \succeq (\mathcal{M}, P)$ be such that $\mathcal{N}$ is strongly κ-homogeneous and $(\mathcal{N}, Q)$ is κ-saturated, where $\kappa > \operatorname{card}(A)$, obtained using Proposition 7.12. Then define $\Phi\colon S_n(T_A) \to [0,1]$ by $\Phi(p) = Q(b)$ for $b \in N^n$ realizing p. We first need to show Φ is well-defined. On the one hand, $\mathcal{N}_A$ realizes every type $p \in S_n(T_A)$ (so b exists). On the other hand, Q is assumed to be $\operatorname{Aut}_A(\mathcal{N})$-invariant and this group acts transitively on the set of realizations of any given $p \in S_n(T_A)$.

Next we show that Φ is continuous with respect to the logic topology on $S_n(T_A)$. Fix $p \in S_n(T_A)$ and let $r = \Phi(p)$. For any $b \in N^n$ we have the implication

$$b \text{ realizes } p \text{ in } \mathcal{N}_A \quad \Rightarrow \quad Q(b) = r.$$

Since $(\mathcal{N}, Q)$ is κ-saturated, it follows that for each $\epsilon > 0$ there exists a condition $\varphi = 0$ in p and $\delta > 0$ so that for any $b \in N^n$ we have the implication

$$\varphi^{\mathcal{N}}(b) < \delta \quad \Rightarrow \quad |Q(b) - r| \leq \epsilon/2.$$

Therefore Φ maps $[\varphi < \delta]$, which is a logic neighborhood of p, into the open interval $(r-\epsilon, r+\epsilon)$. Hence Φ is continuous.

By Theorem 9.9 we conclude that P is definable in $\mathcal{M}$ over A. □

Distance predicates

Let $\mathcal{M}$ be an L-structure and $D \subseteq M^n$. The predicate giving the distance in M^n to D is given by

$$\operatorname{dist}(x, D) = \inf\{d(x,y) \mid y \in D\}.$$

Here $x = (x_1, \ldots, x_n)$ and $y = (y_1, \ldots, y_n)$ range over M^n and we consider the metric d on M^n defined by

$$d(x,y) = \max\big(d(x_1,y_1), \ldots, d(x_n,y_n)\big).$$

Predicates of the form $\operatorname{dist}(x, D)$ are important in the model theory of metric structures. We show next that they can be characterized by axioms in continuous logic.

Consider a predicate $P \colon M^n \to [0,1]$ and the following conditions.

$$(E_1) \qquad \sup_x \inf_y \max\left(P(y), |P(x) - d(x,y)|\right) = 0;$$

$$(E_2) \qquad \sup_x |P(x) - \inf_y \min\left(P(y) + d(x,y), 1\right)| = 0.$$

Observe that for any $D \subseteq M^n$, $P(x) = \operatorname{dist}(x, D)$ satisfies E_1 and E_2.

9.12 Theorem. *Let $(\mathcal{M}, F)$ be an L-structure satisfying conditions E_1 and E_2. Let $D = \{x \in M^n \mid F(x) = 0\}$ be the zeroset of F. Then $F(x) = \operatorname{dist}(x, D)$ for all $x \in M^n$.*

Proof By E_2, $F(x) \le F(y) + d(x,y)$ for all y. So for $y \in D$, $F(x) \le d(x,y)$. So $F(x) \le \operatorname{dist}(x, D)$.

Fix $\epsilon > 0$. We will show that $\operatorname{dist}(x, D) \le F(x) + \epsilon$ for all $x \in M^n$. Letting ϵ go to 0 will complete the proof. We generate a sequence (y_k) in $\mathcal{M}$ using E_1. Set $y_1 = x$, any fixed element of M^n. Choose y_2 such that $F(y_2) \le \frac{\epsilon}{8}$ and $|F(x) - d(x, y_2)| \le \frac{\epsilon}{8}$. Continue by induction, satisfying

$$F(y_{k+1}) \le \frac{\epsilon}{2^{k+2}} \qquad |F(y_k) - d(y_k, y_{k+1})| \le \frac{\epsilon}{2^{k+2}}.$$

Therefore, $d(y_k, y_{k+1}) \le F(y_k) + |F(y_k) - d(y_k, y_{k+1})| \le \epsilon 2^{-k}$. So (y_k)

is a Cauchy sequence, and hence it converges to some $y \in M^n$. By the continuity of F, $F(y) = 0$. Moreover,

$$d(x, y) = \lim_{k\to\infty} d(y_1, y_{k+1}) \leq d(y_1, y_2) + \sum_{k=2}^{\infty} d(y_k, y_{k+1}) \leq F(x) + \epsilon.$$

Since $y \in D$, this shows $\operatorname{dist}(x, D) \leq F(x) + \epsilon$, as desired. □

Zerosets

Next we turn to definability for *sets* (*i.e.*, subsets of M^n). An obvious way to carry over definability for sets from first-order logic would be to regard a condition $\varphi(x) = 0$ as defining the set of x that satisfy it. This is indeed an important kind of definability for sets, but it turns out to be somewhat weak, due to fact that the rules for constructing formulas in continuous logic are rather generous. This is especially true if one puts zerosets of definable predicates on the same footing as zerosets of formulas (as one should).

In this subsection we briefly explore this notion of definability, and in the next subsection we introduce a stronger, less obvious, but very important kind of definability for sets in metric structures.

9.13 Definition. Let $D \subseteq M^n$. We say that D is a *zeroset* in $\mathcal{M}$ over A if there is a predicate $P\colon M^n \to [0, 1]$ definable in $\mathcal{M}$ over A such that $D = \{x \in M^n \mid P(x) = 0\}$.

The next result shows that zerosets are the same as type-definable sets, with the restriction that the partial type is countable.

9.14 Proposition. *For $D \subseteq M^n$, the following are equivalent.*

(1) *D is a zeroset in $\mathcal{M}$.*
(2) *there is a sequence $(\varphi_m \mid m \geq 1)$ of L-formulas such that*

$$\begin{aligned} D &= \left\{x \in M^n \mid \varphi_m^{\mathcal{M}}(x) = 0 \text{ for all } m \in \mathbb{N}\right\} \\ &= \bigcap_{m=1}^{\infty} \text{Zeroset of } \varphi_m^{\mathcal{M}}. \end{aligned}$$

Proof (1 ⇒ 2): Suppose $D = \{x \in M^n \mid P(x) = 0\}$ and $(\varphi_m^{\mathcal{M}} \mid m \geq 1)$ are L-formulas such that for all $x \in M^n$ and all $m \geq 1$

$$|P(x) - \varphi_m^{\mathcal{M}}(x)| \leq \frac{1}{m}.$$

Then $D = \bigcap_m D_m$, where D_m is the zero set of the interpretation of the L-formula $\left(\varphi_m(x) \dot{-} \frac{1}{m}\right)$ in $\mathcal{M}$.

(2 $\Rightarrow$ 1): If $D = \{x \in M^n \mid \varphi_m^{\mathcal{M}}(x) = 0 \text{ for all } m \in \mathbb{N}\}$, then the definable predicate

$$P(x) = \sum_{m=1}^{\infty} 2^{-m} \varphi_m^{\mathcal{M}}(x).$$

has D as its zeroset. □

9.15 Corollary. *The collection of zerosets in $\mathcal{M}$ over A is closed under countable intersections.*

Definability of sets

The next definition gives what we believe is the correct concept of first-order definability for sets in metric structures.

9.16 Definition. A closed set $D \subseteq M^n$ is *definable* in $\mathcal{M}$ over A if and only if the distance predicate $\operatorname{dist}(x, D)$ is definable in $\mathcal{M}$ over A.

The importance of this concept of definability for sets is shown by the following result. It says essentially that in continuous first-order logic we will retain definability of predicates if we quantify (using sup or inf) over sets that are definable in this sense, but not if we quantify over other sets.

9.17 Theorem. *For a closed set $D \subseteq M^n$ the following are equivalent:*
(1) D is definable in $\mathcal{M}$ over A.
(2) For any predicate $P\colon M^m \times M^n \to [0,1]$ that is definable in $\mathcal{M}$ over A, the predicate $Q\colon M^m \to [0,1]$ defined by

$$Q(x) = \inf\{P(x,y) \mid y \in D\}$$

is definable in $\mathcal{M}$ over A.

Proof To prove (1) we only need to assume (2) for the case in which $P(x,y) = d(x,y)$ (so P is the interpretation of a quantifier-free formula); here $m = n$.

So assume D is definable in $\mathcal{M}$ over A. Let $P\colon M^m \times M^n \to [0,1]$ be any predicate that is definable in $\mathcal{M}$ over A. This ensures that P is uniformly continuous, so (see Proposition 2.10) there is an increasing, continuous function $\alpha\colon [0,1] \to [0,1]$ with $\alpha(0) = 0$ such that

$$|P(x,y) - P(x,z)| \leq \alpha(d(y,z))$$

for any $x \in M^m$ and $y, z \in M^n$. Let Q be defined on M^m by $Q(x) = \inf\{P(x,y) \mid y \in D\}$.

We will show that $Q(x) = \inf\{P(x,z) + \alpha(\operatorname{dist}(z,D)) \mid z \in M^n\}$ for all $x \in M^m$. This shows that Q is definable in $\mathcal{M}$ over A, given that P and $\operatorname{dist}(z,D)$ are definable in $\mathcal{M}$ over A. (Notice that we replaced the inf over D by the inf over M^n and it is expressible by one of the basic constructs of continuous logic.)

By our choice of α we have $P(x,y) \leq P(x,z) + \alpha(d(y,z))$ for all $x \in M^m$ and $y, z \in M^n$. Taking the inf over $y \in D$ and using the fact that α is continuous and increasing, we conclude that

$$Q(x) \leq P(x,z) + \alpha(\operatorname{dist}(z,D))$$

for all $x \in M^m$ and $z \in M^n$. Taking the inf of the right side over $z \in D$ yields $Q(x)$, and the inf over $z \in M^n$ is even smaller, but is bounded below by $Q(x)$. This shows

$$Q(x) = \inf\{P(x,z) + \alpha(\operatorname{dist}(z,D)) \mid z \in M^n\}$$

for all $x \in M^m$ and completes the proof of (1) ⇒ (2). □

The following result shows another useful property of definable sets that need not be true of zerosets.

9.18 Proposition. *Let $\mathcal{N} \preceq \mathcal{M}$ be L-structures, and let $D \subseteq M^n$ be definable in $\mathcal{M}$ over A, where $A \subseteq N$. Then:*
(1) For any $x \in N^n$, $\operatorname{dist}(x,D) = \operatorname{dist}(x, D \cap N^n)$. Thus $D \cap N^n$ is definable in $\mathcal{N}$ over A.
(2) $(\mathcal{N}, \operatorname{dist}(\cdot, D \cap N^n)) \preceq (\mathcal{M}, \operatorname{dist}(\cdot, D))$.
(3) If $D \neq \emptyset$, then $D \cap N^n \neq \emptyset$.

Proof Statement (1) is an immediate consequence of Theorem 9.12 and (2) follows from (1) by Proposition 9.7. For (3), note that if $D \neq \emptyset$, then $\inf_x \operatorname{dist}(x,D) = 0$ in $\mathcal{M}$. Therefore $\inf_x \operatorname{dist}(x, D \cap N^n) = 0$ in $\mathcal{N}$ by (2). Hence there exists $a \in N$ such that $\operatorname{dist}(a, D \cap N^n) < 1$. This implies $D \cap N^n \neq \emptyset$, since otherwise $\operatorname{dist}(a, D \cap N^n) = 1$. □

If $D \subseteq M^n$ is definable in $\mathcal{M}$, then it is certainly a zeroset (it is the set of zeros of the definable predicate $\operatorname{dist}(x,D)$). However, the converse is not generally true. The next result explores the distinction between these two concepts.

9.19 Proposition. *For a closed set $D \subseteq M^n$, the following are equivalent:*

(1) *D is definable in $\mathcal{M}$ over A.*
(2) *There is a predicate $P\colon M^n \to [0,1]$, definable in $\mathcal{M}$ over A, such that $P(x) = 0$ for all $x \in D$ and*

$$\forall \epsilon > 0 \ \exists \delta > 0 \ \forall x \in M^n \ \ (P(x) \le \delta \Rightarrow \operatorname{dist}(x, D) \le \epsilon) .$$

(3) *There is a sequence $(\varphi_m \mid m \ge 1)$ of $L(A)$-formulas and a sequence $(\delta_m \mid m \ge 1)$ of positive real numbers such that for all $m \ge 1$ and $x \in M^n$,*

$$\big(x \in D \Rightarrow \varphi_m^{\mathcal{M}}(x) = 0\big) \ \textit{and}$$

$$\left(\varphi_m^{\mathcal{M}}(x) \le \delta_m \Rightarrow \operatorname{dist}(x, D) \le \frac{1}{m}\right).$$

Proof $(1 \Rightarrow 3)$: Let $F(x) = \operatorname{dist}(x, D)$ and assume it is definable in $\mathcal{M}$ over A. So there exists a sequence $(\psi_m(x) \mid m \ge 1)$ of $L(A)$-formulas such that for all $x \in M^n$ and $m \ge 1$ we have

$$|F(x) - \psi_m^{\mathcal{M}}(x)| \le \frac{1}{3m}.$$

If $x \in D$, then $F(x) = 0$ and so $\psi_m^{\mathcal{M}}(x) \le \frac{1}{3m}$. Also, if $\psi_m^{\mathcal{M}}(x) \le \frac{2}{3m}$ we have

$$F(x) \le \psi_m^{\mathcal{M}}(x) + |F(x) - \psi_m^{\mathcal{M}}(x)| \le \frac{2}{3m} + \frac{1}{3m} = \frac{1}{m}.$$

Hence the $L(A)$-formulas $\varphi_m(x) = \big(\psi_m^{\mathcal{M}}(x) \dot{-} \frac{1}{3m}\big)$ have the desired property (with $\delta_m = \frac{2}{3m}$.)

$(3 \Rightarrow 2)$: Set $P(x) = \sum_{m=1}^{\infty} 2^{-m} \varphi_m^{\mathcal{M}}(x)$.

$(2 \Rightarrow 1)$: We use Proposition 2.10. This gives us a continuous, increasing function $\alpha\colon [0,1] \to [0,1]$ such that $\alpha(0) = 0$ and for all $x \in M^n$, $\operatorname{dist}(x, D) \le \alpha(P(x))$.

Consider the function

$$F(x) = \inf_y \min\left(\alpha\left(P(y)\right) + d(x, y), 1\right).$$

First of all, this predicate is definable in $\mathcal{M}$ over A, since P is definable in $\mathcal{M}$ over A. Indeed, if $(\varphi_n \mid n \geq 1)$ are $L(A)$-formulas such that $\varphi_n^{\mathcal{M}}$ converges to P on M^n, then

$$\psi_n = \inf_y \min\left(\alpha(\varphi_n(y)) + d(x,y), 1\right)$$

gives a sequence of $L(A)$-formulas such that $\psi_n^{\mathcal{M}}$ converges uniformly to F.

Second, we observe that $F(x) = \mathrm{dist}(x, D)$ for all $x \in M^n$. If y is any element of D (so $P(y) = 0$), we see that

$$F(x) \leq \min\left(\alpha(0) + d(x,y), 1\right) = d(x,y)$$

and hence $F(x) \leq \mathrm{dist}(x, D)$.

On the other hand, we have $\alpha\left(P(y)\right) \geq \mathrm{dist}(y, D)$ for all y, and hence

$$\begin{aligned} F(x) &\geq \inf_y \min\left(\mathrm{dist}(y, D) + d(x,y), 1\right) \\ &\geq \min\left(\mathrm{dist}(x, D), 1\right). \end{aligned}$$

Since the metric d is bounded by 1, this shows $F(x) \geq \mathrm{dist}(x, D)$, as desired. □

9.20 Remark. Suppose $\mathcal{M}$ is an ω_1-saturated L-structure and A is a countable subset of M.

If $P, Q\colon M^n \to [0,1]$ are definable in $\mathcal{M}$ over A and P, Q have the same zeroset, then an easy saturation argument shows that

$$\forall \epsilon > 0\ \exists \delta > 0\ \forall x \in M^n\ \ (P(x) \leq \delta \Rightarrow Q(x) \leq \epsilon)$$

and

$$\forall \epsilon > 0\ \exists \delta > 0\ \forall x \in M^n\ \ (Q(x) \leq \delta \Rightarrow P(x) \leq \epsilon).$$

Now suppose $P\colon M^n \to [0,1]$ is definable in $\mathcal{M}$ over A. The proof of $(2 \Rightarrow 1)$ in Proposition 9.19 together with the observation in the previous paragraph yields the following:

The zeroset D of P is definable in $\mathcal{M}$ over A if and only if P satisfies the statement in Proposition 9.19(2):

$$\forall \epsilon > 0\ \exists \delta > 0\ \forall x \in M^n\ \ (P(x) \leq \delta \Rightarrow \mathrm{dist}(x, D) \leq \epsilon).$$

As an example, we use this to give a geometric criterion for definability of closed balls $B(a,r) = \{x \in M^n \mid d(a,x) \leq r\}$; namely, $B(a,r)$ is definable in $\mathcal{M}$ over $\{a\}$ if and only if the family $(B(a, r+\delta) \mid \delta > 0)$ of closed balls converges to $B(a,r)$, in the sense of the Hausdorff metric, as $\delta \to 0$. (To prove this, apply the preceding statement to the predicate

$P(x) = (d(a,x) \dot{-} r)$, which is definable in $\mathcal{M}$ over $\{a\}$ and has $B(a,r)$ as its zeroset. Note that $P(x) \leq \delta$ iff $d(a,x) \leq r + \delta$.)

9.21 Remark. Let $\mathcal{M}$ be a discrete structure from ordinary first-order logic. We explore the meaning of these notions of definability for sets in $\mathcal{M}$, regarding it as a metric structure in which we take the distance between distinct elements always to be 1.

Note that in the metric setting there are more formulas than the usual first-order ones. For example, $\psi = \sum_{k=1}^{N} 2^{-k}\varphi_k$ is a formula whenever $(\varphi_k \mid 1 \leq k \leq N)$ are first-order formulas; for $x \in M^n$, the possible values of $\psi^{\mathcal{M}}$ are $\frac{k}{2^N}$, where $k = 0, \ldots, 2^N - 1$.

What one can easily prove here, by induction on formulas, is the following: for any L-formula φ, $\{\varphi^{\mathcal{M}}(x) \mid x \in M^n\}$ is a finite set. Moreover, for any $r \in [0,1]$, $\{x \in M^n \mid \varphi^{\mathcal{M}}(x) = r\}$ is definable in $\mathcal{M}$ by an ordinary first-order formula.

For any $D \subseteq M^n$ one has the following characterizations:

(1) D is definable in $\mathcal{M}$ over A if and only if D is definable in $\mathcal{M}$ over A by a first-order formula, in the usual sense of model theory.
(2) D is a zeroset in $\mathcal{M}$ over A if and only if D is the intersection of countably many sets definable in $\mathcal{M}$ over A by first-order formulas.

To prove (1) above, suppose $\mathrm{dist}(x,D)$ is definable in $\mathcal{M}$ over A. In this setting, $\mathrm{dist}(x,D)$ is simply $1 - \chi_D$, where χ_D denotes the characteristic function of D. Let φ be an $L(A)$-formula such that for all $x \in M^n$

$$|\mathrm{dist}(x,D) - \varphi^{\mathcal{M}}(x)| \leq \frac{1}{3}.$$

Then $D = \left\{x \in M^n \mid \varphi^{\mathcal{M}}(x) \leq \frac{1}{2}\right\}$ and this is first-order definable in $\mathcal{M}$ over A.

Definability of functions

We give a brief introduction to definability of functions (from M^n into M) in a metric structure $\mathcal{M}$.

9.22 Definition. Let $\mathcal{M}$ be an L-structure and $A \subseteq M$ and consider a function $f\colon M^n \to M$. We say f is definable in $\mathcal{M}$ over A if and only if the function $d(f(x),y)$ on M^{n+1} is a predicate definable in $\mathcal{M}$ over A.

9.23 Proposition. *If the function $f\colon M^n \to M$ is definable in $\mathcal{M}$ over A, then f is uniformly continuous; indeed, any modulus of uniform continuity for the predicate $d(f(x), y)$ is a modulus of uniform continuity for f.*

Proof Suppose $\Delta\colon (0,1] \to (0,1]$ is a modulus of uniform continuity for $d(f(x), y)$. That is, for any $x, x' \in M^n$ and $y, y' \in M$ and for any $\epsilon \in (0,1]$ we know that if $d(x, x') < \Delta(\epsilon)$ and $d(y, y') < \Delta(\epsilon)$, then $|d(f(x), y) - d(f(x'), y')| \le \epsilon$. Taking $y' = y = f(x')$ we get that $d(x, x') < \Delta(\epsilon)$ implies $d(f(x), f(x')) \le \epsilon$. □

For any function $f\colon M^n \to M$ we denote the graph of f by $\mathcal{G}_f$. We regard $\mathcal{G}_f$ as a subset of M^{n+1}.

Note that if f is definable in $\mathcal{M}$ over A then its graph $\mathcal{G}_f$ is definable in $\mathcal{M}$ over A as a subset of M^{n+1}. This follows from the identity

$$\operatorname{dist}\big((x, y), \mathcal{G}_f\big) = \inf_z \max\big(d(x, z), d(f(z), y)\big)$$

(in which x, z range over M^n and y ranges over M).

The converse of this observation is true in a strong form, if we work in a sufficiently saturated model.

9.24 Proposition. *Let $\mathcal{M}$ be κ-saturated, where κ is uncountable, and let $A \subseteq M$ have cardinality $< \kappa$. Let $f\colon M^n \to M$ be any function. The following are equivalent:*
(1) f is definable in $\mathcal{M}$ over A.
(2) $\mathcal{G}_f$ is type-definable in $\mathcal{M}$ over A.

Proof It remains to prove (2) $\Rightarrow$ (1). Write $P(x, y) = d(f(x), y)$ for $x \in M^n$ and $y \in M$. Let $\Gamma(x, y)$ be a set of $L(A)$-conditions that type-defines $\mathcal{G}_f$ in $\mathcal{M}$. Fix $r \in [0,1]$ and note that

$$P(x, y) \le r \Leftrightarrow \exists z\big((x, z) \in \mathcal{G}_f \wedge d(z, y) \le r\big).$$

This shows that the set $\{(x, y) \in M^{n+1} \mid P(x, y) \le r\}$ is type-defined in $\mathcal{M}$ by the set of $L(A)$-conditions of the form

$$\inf_z \max\big(\varphi(x, z), d(z, y) \dot{-} r\big) = 0$$

where $\varphi = 0$ is any condition in Γ. A similar argument shows that the set $\{(x, y) \in M^{n+1} \mid P(x, y) \ge r\}$ is type-definable in $\mathcal{M}$ over A for each $r \in [0,1]$. By Corollary 9.10 this shows that P is definable and hence that f is definable in $\mathcal{M}$ over A. □

9.25 Proposition. *Let $\mathcal{M}$ be an L-structure and $A \subseteq M$. Suppose the function $f\colon M^n \to M$ is definable in $\mathcal{M}$ over A. Then:*
(1) If $\mathcal{N} \preceq \mathcal{M}$ and $A \subseteq N$, then f maps N^n into N and the restriction of f to N^n is definable in $\mathcal{N}$ over A.
(2) If $\mathcal{N} \succeq \mathcal{M}$ then there is a function $g\colon N^n \to N$ such that g extends f and g is definable in $\mathcal{N}$ over A.

Proof (1) Fix any $(a_1, \dots, a_n) \in N^n$ and let $b = f(a_1, \dots, a_n) \in M$. Let $P\colon M \to [0,1]$ be the predicate defined by $P(y) = d(b, y)$ for all $y \in M$; note that P is definable in $\mathcal{M}$ over $A \cup \{a_1, \dots, a_n\} \subseteq N$. Let $Q\colon N \to [0,1]$ be the restriction of P to N; evidently Q is definable in $\mathcal{N}$ over $A \cup \{a_1, \dots, a_n\}$ and, by Proposition 9.7, $(\mathcal{N}, Q) \preceq (\mathcal{M}, P)$. Note that P satisfies $\inf_y P(y) = 0$ and $d(x, y) \leq P(x) + P(y)$, where x, y range over M. Transferring these conditions from $(\mathcal{M}, P)$ to $(\mathcal{N}, Q)$ we may find a sequence $(c_k \mid k \geq 1)$ in N satisfying $Q(c_k) \leq 1/k$ for all $k \geq 1$ and hence also $d(c_k, c_l) \leq 1/k + 1/l$ for all $k, l \geq 1$. It follows that $(c_k \mid k \geq 1)$ converges to an element of N which is a zero of Q and hence a zero of P. This limit must be b, since b is the only zero of P.
(2) By making $\mathcal{N}$ larger we may assume that it is ω_1-saturated. Using (1) it suffices to prove (2) for the larger elementary extension. Let $P\colon M^{n+1} \to [0,1]$ be the predicate $P(x, y) = d(f(x), y)$. Using Proposition 9.8 let $Q\colon N^{n+1} \to [0,1]$ be the predicate that extends P and is definable in $\mathcal{N}$ over A, so we have $(\mathcal{N}, Q) \succeq (\mathcal{M}, P)$. Note that P satisfies $\sup_x \inf_y P(x, y) = 0$. Hence the same is true of Q. Using this and the fact that $\mathcal{N}$ is ω_1-saturated it follows that for all $x \in N^n$ there exists at least one $y \in N$ such that $Q(x, y) = 0$. Note also that from the definition of P and the triangle inequality for d it follows that P satisfies

$$\sup_x \sup_y \sup_{y'} \big(|d(y', y) - P(x, y')| \dot{-} P(x, y) \big) = 0.$$

Hence the same is true of Q. From this it follows that for each $x \in N^n$ there is at most one $y \in N$ such that $Q(x, y) = 0$. Therefore the zero set of Q is the graph of some function $g\colon N^n \to N$. Moreover, it also follows that if $Q(x, y) = 0$ then $Q(x, y') = d(y, y')$ for all $y' \in N$. That is, for all $x \in N^n$ and $y' \in N$ we have $Q(x, y') = d(g(x), y')$. This shows that the function g is definable in $\mathcal{N}$ over A as desired. □

The preceding results permit one to show easily that the collection of functions definable in a given structure is closed under composition and that the result of substituting definable functions into a definable predicate is a definable predicate.

Extension by definition

Here we broaden our perspective to consider our definability concepts in a *uniform* way, relative to the class of all models of a satisfiable L-theory T. The earlier parts of this section can be interpreted as dealing with the case where T is *complete*; now we consider more general theories.

Let L be any signature for metric structures and let T be any satisfiable L-theory; let L_0 be a signature contained in L.

If T_0 is an L_0-theory, we say that T_0 is the *restriction of* T *to* L_0 (or, equivalently, that T is a *conservative extension of* T_0) if for every closed L_0-condition E we have

$$T \models E \quad \text{iff} \quad T_0 \models E.$$

Note that for T to be a conservative extension of T_0 it suffices (but need not be necessary) to require that T is an extension of T_0 and that every model of T_0 has an expansion that is a model of T.

9.26 Definition. An L-formula $\varphi(x_1, \dots, x_n)$ is *defined in* T *over* L_0 if for each $\epsilon > 0$ there exists an L_0-formula $\psi(x_1, \dots, x_n)$ such that

$$T \models \big(\sup_{x_1} \dots \sup_{x_n} |\varphi - \psi|\big) \leq \epsilon.$$

If P is an n-ary predicate symbol in L, we say P is *defined in* T *over* L_0 if the formula $P(x_1, \dots, x_n)$ is defined in T over L_0. If f is an n-ary function symbol in L, we say f is *defined in* T *over* L_0 if the formula $d(f(x_1, \dots, x_n), y)$ is defined in T over L_0. In particular, if c is a constant symbol in L, we say c is *defined in* T *over* L_0 if the formula $d(c, y)$ is defined in T over L_0.

9.27 Definition. Let T, L, L_0 be as above and let T_0 be an L_0-theory. We say T is an *extension by definitions of* T_0 if T is a conservative extension of T_0 and every nonlogical symbol in L is defined in T over L_0.

Let T, L, L_0 be as above and let T_0 be the restriction of T to L_0. For each $n \geq 0$ define $\pi_n \colon S_n(T) \to S_n(T_0)$ to be the restriction map: $\pi_n(p)$ is the set of all L_0-conditions in p, for any $p \in S_n(T)$. Evidently each π_n is continuous with respect to the logic topologies. Using Theorem 8.12(1,2) we see that π_n is surjective for each n. (Since $\pi_n(S_n(T))$ is a closed subset of $S_n(T_0)$, there is a set $\Gamma(x_1, \dots, x_n)$ of L_0-conditions such that $\pi_n(S_n(T))$ is the set of all $q \in S_n(T_0)$ such that $\Gamma \subseteq q$. If

$\varphi(x_1, \ldots, x_n) = 0$ is any condition in Γ we have

$$T \models (\sup_{x_1} \ldots \sup_{x_n} \varphi(x_1, \ldots, x_n)) = 0$$

and therefore

$$T_0 \models (\sup_{x_1} \ldots \sup_{x_n} \varphi(x_1, \ldots, x_n)) = 0.$$

It follows that every $q \in S_n(T_0)$ contains Γ.)

9.28 Proposition. *Let T, L, L_0 be as above and suppose T is an extension by definitions of the L_0-theory T_0. Then every L-formula is defined in T over L_0.*

Proof Note that every model $\mathcal{M}$ of T is completely determined by its reduct $\mathcal{M}|L_0$ to L_0. We will use this to show that the restriction map π_n is injective for each n. Therefore π_n^{-1} is a homeomorphism from $S_n(T_0)$ onto $S_n(T)$ for each n. Using Theorem 8.12(3) we conclude that every L-formula is defined in T over L_0.

To complete the proof, suppose $p_1, p_2 \in S_n(T)$ have $\pi_n(p_1) = \pi_n(p_2)$. For each $j = 1, 2$ let $\mathcal{M}_j$ be a model of T and a_j an n-tuple in M_j that realizes p_j in $\mathcal{M}_j$. It follows that $(\mathcal{M}_1|L_0, a_1)$ and $(\mathcal{M}_2|L_0, a_2)$ are elementarily equivalent. By Theorem 5.7 there is an ultrafilter D such that the ultrapowers $(\mathcal{M}_1|L_0, a_1)_D$ and $(\mathcal{M}_2|L_0, a_2)_D$ are isomorphic, say by the function f from $(\mathcal{M}_1|L_0)_D$ onto $(\mathcal{M}_2|L_0)_D$. Let $\mathcal{N}$ be the unique L-structure for which f is an isomorphism from $(\mathcal{M}_1)_D$ onto $\mathcal{N}$. This ensures that $\mathcal{N}$ is a model of T and $\mathcal{N}|L_0 = (\mathcal{M}_2|L_0)_D = (\mathcal{M}_2)_D|L_0$. Therefore $\mathcal{N}$ and $(\mathcal{M}_2)_D$ are identical, and hence f is an isomorphism from $(\mathcal{M}_1)_D$ onto $(\mathcal{M}_2)_D$. Since f maps a_1 onto a_2, this shows $p_1 = p_2$, as desired. □

9.29 Corollary. *The property of being an extension by definitions is transitive. That is, if T_1 is an extension by definitions of T_0 and T_2 is an extension by definitions of T_1, then T_2 is an extension by definitions of T_0.*

Proof Immediate from the previous result. □

9.30 Remark. The following observation is useful when constructing extensions by definition in the metric setting. Suppose γ is an ordinal and $(T_\alpha \mid \alpha < \gamma)$ are theories in continuous logic such that (a) $T_{\alpha+1}$

is an extension by definitions of T_α whenever $\alpha + 1 < \gamma$ and (b) $T_\lambda = \cup(T_\alpha \mid \alpha < \lambda)$ whenever λ is a limit ordinal $< \gamma$. Let $T = \cup(T_\alpha \mid \alpha < \gamma)$. Then T is an extension by definitions of T_0. (The proof is by induction on γ; Corollary 9.29 takes care of the case when γ is a successor ordinal and the limit ordinal case is trivial.)

9.31 Corollary. *Let T, L, L_0 be as above and let T_0 be an L_0-theory. If T is an extension by definitions of T_0, then every model of T_0 has a (unique) expansion to a model of T.*

Proof We have already noted the uniqueness of the expansion, so only its existence needs to be proved. Let $\mathcal{M}_0$ be any model of T_0. As shown in the proof of Proposition 9.28, there is a unique complete L-theory T_1 that contains T and $\mathrm{Th}(\mathcal{M}_0)$. ($T_1$ is the unique element of $S_0(T)$ that satisfies $\pi_0(T_1) = \mathrm{Th}(\mathcal{M}_0)$.) Let $\mathcal{M}$ be a $\mathrm{card}(\mathcal{M}_0)^+$-saturated model of T_1. Since $\mathcal{M}|L_0 \equiv \mathcal{M}_0$, we may assume without loss of generality that $\mathcal{M}_0 \preceq \mathcal{M}|L_0$. A simple modification of the proof of Proposition 9.25(1) shows that the universe of $\mathcal{M}_0$ is closed under $f^{\mathcal{M}}$ for every function symbol f of L and contains $c^{\mathcal{M}}$ for every constant symbol c of L. Hence there exists a substructure $\mathcal{M}_0^*$ of $\mathcal{M}$ whose reduct to L_0 equals $\mathcal{M}_0$. It follows from Proposition 9.28 and the fact that $\mathcal{M}_0 \preceq \mathcal{M}|L_0$ that $\mathcal{M}_0^*$ is an elementary substructure of $\mathcal{M}$. In particular, $\mathcal{M}_0^*$ is a model of T and an expansion of $\mathcal{M}_0$, as desired. □

Now we describe certain standard ways of obtaining extensions by definition of a given L_0-theory T_0. Typically this is done via a sequence of steps, in each of which we add a definable predicate, or a definable constant, or a definable function. As the basis of each step we have in hand a previously constructed theory T that is an extension by definitions of T_0, with L being the signature of T.

First, consider the case where we want to add a definable n-ary predicate. Here we have a sequence $(\varphi_k(x_1, \dots, x_n) \mid k \geq 1)$ of L-formulas that is uniformly Cauchy in all models of T, in the sense that the following statement holds:

$$\forall \epsilon > 0 \; \exists N \; \forall k, l > N \; T \models \big(\sup_{x_1} \dots \sup_{x_n} |\varphi_k - \varphi_l|\big) \leq \epsilon.$$

Choose an increasing sequence of positive integers $(N_m \mid m \geq 1)$ such that

$$\forall m \geq 1 \; \forall k, l > N_m \; T \models \big(\sup_{x_1} \dots \sup_{x_n} |\varphi_k - \varphi_l|\big) \leq 2^{-m}.$$

We then let P be a new n-ary predicate symbol and take T' to be the $L(P)$-theory obtained by adding to T the conditions

$$\big(\sup_{x_1} \dots \sup_{x_n} |\varphi_{k(m)}(x_1, \dots, x_n) - P(x_1, \dots, x_n)|\big) \leq 2^{-m}$$

for every $m \geq 1$ and $k(m) = N_m + 1$. Then every model of T has an expansion that is a model of T', and P is defined in T' over L, by construction. Therefore T' is an extension by definitions of T. Hence T' is an extension by definitions of T_0 by Corollary 9.29. Given any model $\mathcal{M}$ of T', this implies that the predicate $P^{\mathcal{M}}$ is definable (over $\emptyset$) in the reduct $\mathcal{M}|L_0$. More precisely, $P^{\mathcal{M}}$ is the uniform limit of the predicates $\varphi_m^{\mathcal{M}|L}$ as $m \to \infty$, and each of these predicates is definable in $\mathcal{M}|L_0$.

Note that in defining the signature $L(P)$ we need to specify a modulus of uniform continuity for the predicate symbol P. Such a modulus can be defined from the sequence $(N_m \mid m \geq 1)$ together with moduli for the formulas $(\varphi_k \mid k \geq 1)$ as indicated in the proof of Proposition 2.5.

Next consider the case where we want to add a definable constant. Without loss of generality we may assume that we have an L-formula $\varphi(y)$ such that

$$T \models \big(\inf_z \sup_y |d(z, y) - \varphi(y)|\big) = 0.$$

This implies that in every model $\mathcal{M}$ of T the zeroset of $\varphi^{\mathcal{M}}$ has a single element, by Proposition 9.25(1).

We then let c be a new constant symbol and take T' to be the $L(c)$-theory obtained by adding to T the condition $(\sup_y |d(c, y) - \varphi(y)|) = 0$. Again we have that every model of T has an expansion that is a model of T' and c is definable in T' over L. Hence T' is an extension by definitions of T_0.

Finally, generalizing the case of adding a definable constant, consider the case where we want to add a definable function. Without loss of generality we may assume that we have an L-formula $\varphi(x_1, \dots, x_n, y)$ such that

$$T \models \big(\sup_{x_1} \dots \sup_{x_n} \inf_z \sup_y |d(z, y) - \varphi(x_1, \dots, x_n, y)|\big) = 0.$$

This implies that in every model $\mathcal{M}$ of T the zeroset of $\varphi^{\mathcal{M}}$ is the graph of a total function from M^n to M, by Proposition 9.25(1).

We then let f be a new n-ary function symbol and take T' to be the $L(f)$-theory obtained by adding to T the condition

$$\big(\sup_{x_1} \dots \sup_{x_n} \sup_y |d(f(x_1, \dots, x_n), y) - \varphi(x_1, \dots, x_n, y)|\big) = 0.$$

Again we have that every model of T has an expansion that is a model of T' and f is definable in T' over L. Hence T' is an extension by definitions of T_0. Note that when introducing the signature $L(f)$ we must specify a modulus of uniform continuity for f. This can be taken to be the modulus of uniform continuity of the L-formula φ. (See the proof of Proposition 9.23.)

In order to simplify matters, we have described the addition of a constant c or function f only in the apparently restricted situation where the definitions of $d(c, y)$ or $d(f(x_1, \ldots, x_n), y)$ are given by *formulas* of L rather than by definable predicates. We want to emphasize that this is not a real limitation; it would be overcome by first adding the needed definable predicate and then using it to add the desired constant or function.

We conclude this section with a result that generalizes Beth's Definability Theorem to continuous logic.

9.32 Theorem. *Let T be an L-theory and let L_0 be a signature contained in L. Let S be any nonlogical symbol in L. Assume that S is implicitly defined in T over L_0; that is, assume that if $\mathcal{M}, \mathcal{N}$ are models of T for which $\mathcal{M}|L_0 = \mathcal{N}|L_0$, then one always has $S^{\mathcal{M}} = S^{\mathcal{N}}$. Then S is defined in T over L_0.*

Proof Let φ be the $L_0(S)$-formula $S(x_1, \ldots, x_n)$ if S is an n-ary predicate symbol, $d(S(x_1, \ldots, x_n), x_{n+1})$ if S is an n-ary function symbol, and $d(S, x_1)$ if S is a constant symbol. Denote the list of variables in φ simply as x. To show that S is defined in T over L_0 we need to show that the formula $\varphi(x)$ is defined in T over L_0.

Let T_1 be the restriction of T to the signature $L_0(S)$. Our first step is to prove that T_1 implicitly defines S over L_0. To prove this we will show that for each $\epsilon > 0$ there exist finitely many L_0-formulas, $\psi_1(x), \ldots, \psi_k(x)$ such that

$$T \models \big(\min_{1 \leq j \leq k} (\sup_x |\varphi(x) - \psi_j(x)|)\big) \leq \epsilon.$$

Suppose this fails for some specific $\epsilon > 0$. Then there exists a model $\mathcal{M}$ of T such that for every L_0-formula $\psi(x)$ we have

$$\mathcal{M} \models (\sup_x |\varphi(x) - \psi_j(x)|) \geq \epsilon.$$

This implies that the predicate $Q = \varphi^{\mathcal{M}}$ is not definable (over $\emptyset$) in $\mathcal{M}|L_0$. By Proposition 7.12 we may assume that $\mathcal{M}$ is ω-saturated and that $\mathcal{M}|L_0$ is strongly ω-homogeneous. As shown in the proof of Corollary 9.11, there is an automorphism τ of $\mathcal{M}|L_0$ such that $Q \neq Q \circ \tau$. Let $\mathcal{N}$ be the unique L-structure for which τ is an isomorphism from $\mathcal{M}$ onto $\mathcal{N}$. Evidently $\mathcal{N} \models T$. Since τ is an automorphism of $\mathcal{M}|L_0$, we have $\mathcal{N}|L_0 = \mathcal{M}|L_0$. However, $\varphi^{\mathcal{N}} = Q \circ \tau \neq Q = \varphi^{\mathcal{M}}$. This contradicts the assumption that S is implicitly defined in T over L_0.

So, we now have that S is implicitly defined in T_1 over L_0. Let T_0 be the restriction of T_1 to L_0 (which is the same as the restriction of T to L_0). For each $n \geq 0$, let $\pi_n \colon S_n(T_1) \to S_n(T_0)$ be the restriction map: $\pi_n(p)$ is equal to the set of all L_0-conditions in p. Arguing as in the proof of Proposition 9.28 we see that each π_n is a surjective homeomorphism and we conclude that the formula $\varphi(x)$ is defined in T_1 over L_0. Hence $\varphi(x)$ is defined in T over L_0, as desired. □

9.33 Corollary. *Let T be an L-theory and T_0 an L_0-theory, both satisfiable, where L_0 is contained in L. Then T is an extension by definitions of T_0 if and only if every model of T_0 has a unique expansion that is a model of T.*

Proof This is immediate from Corollary 9.31 (for the left to right direction) and Theorem 9.32 (for the other direction). □

10 Algebraic and definable closures

In this section we introduce the concepts of *definable* and *algebraic* closure of a set in a metric structure. There are several reasonable choices for the definitions, but they turn out to be equivalent.

10.1 Definition. Let $\mathcal{M}$ be an L-structure and A a subset of M, and let $a \in M^n$. We say that a is *definable* in $\mathcal{M}$ over A if the set $\{a\}$ is definable in $\mathcal{M}$ over A (that is, if the predicate $d(\cdot, a)$ is definable in $\mathcal{M}$ over A). We say that a is *algebraic* in $\mathcal{M}$ over A if there exists a compact set $C \subseteq M^n$ such that $a \in C$ and C is definable in $\mathcal{M}$ over A.

As in the usual first-order setting, these properties of tuples reduce to the corresponding properties of their coordinates:

10.2 Proposition. *Let $\mathcal{M}$ be an L-structure and A a subset of M. Let $a = (a_1, \ldots, a_n) \in M^n$. Then a is definable (resp., algebraic) in $\mathcal{M}$ over A if and only if a_j is definable (resp., algebraic) in $\mathcal{M}$ over A for each $j = 1, \ldots, n$.*

Proof We treat the algebraic case. To prove the left to right direction it suffices to prove that if $C \subseteq M^n$ is compact and definable in $\mathcal{M}$ over A, then its projection C_i onto the i^{th} coordinates is definable for each $i = 1, \dots, n$. (It is obviously compact.) For this we first note that for each i the predicate P_i on M^n defined by

$$P_i(x_1, \dots, x_n) = \inf\{d(x_i, y_i) \mid (y_1, \dots, y_n) \in C\}$$

is definable in $\mathcal{M}$ over A, by Theorem 9.17. Then we have $\mathrm{dist}(x_i, C_i) = P_i(x_i, \dots, x_i)$, showing that C_i is definable in $\mathcal{M}$ over A for each i.

For the right to left direction, suppose $C_i \subseteq M$ is a compact set witnessing that a_i is algebraic in $\mathcal{M}$ over A. Then the product $C = C_1 \times \cdots \times C_n$ witnesses the same property for a. Note that the distance function for such a product is given by

$$\mathrm{dist}\big((x_1, \dots, x_n), C\big) = \inf_{y_1 \in C_1} \cdots \inf_{y_n \in C_n} \max\big(d(x_1, y_1), \dots, d(y_n, x_n)\big)$$

and the right side is definable by n uses of Theorem 9.17. □

Notation 10.3 We let $\mathrm{dcl}_{\mathcal{M}}(A)$ denote the set of all elements of M that are definable in $\mathcal{M}$ over A and we call it the *definable closure* of A in $\mathcal{M}$. Similarly, we let $\mathrm{acl}_{\mathcal{M}}(A)$ denote the set of all elements of M that are algebraic in $\mathcal{M}$ over A and call it the *algebraic closure* of A in $\mathcal{M}$.

We first want to show that the definable and algebraic closures depend only on A and not on the structure in which they are defined. For that we apply some definability results from the previous section.

10.4 Proposition. *Let $\mathcal{M} \preceq \mathcal{N}$ and $A \subseteq M$. Suppose $C \subseteq N^n$ is definable in $\mathcal{N}$ over A and that $C \cap M^n$ is compact. Then C is contained in M^n.*

Proof Let $Q\colon N^n \to [0, 1]$ be the predicate given by $Q(x) = \mathrm{dist}(x, C)$, and assume that Q is definable in $\mathcal{N}$ over A. Let P be the restriction of Q to M^n. By Proposition 9.18 we have that $(\mathcal{M}, P) \preceq (\mathcal{N}, Q)$ and that $P(x) = \mathrm{dist}(x, C \cap M^n)$ for all $x \in M^n$. Fix $\epsilon > 0$ and suppose $c_1, \dots, c_m$ is an ϵ-net in $C \cap M^n$. It follows that whenever $P(x) < \epsilon$ we have $d(x, c_j) \le 2\epsilon$ for some $j = 1, \dots, m$. In other words, the condition

$$\sup_x \min\big(\epsilon \dot- P(x), \min(d(x, c_1), \dots, d(x, c_m)) \dot- 2\epsilon\big) = 0$$

holds in $(\mathcal{M}, P)$. Hence the condition

$$\sup_x \min\left(\epsilon \dot{-} Q(x), \min(d(x, c_1), \ldots, d(x, c_m)) \dot{-} 2\epsilon\right) = 0$$

holds in $(\mathcal{N}, Q)$. It follows that $c_1, \ldots, c_m$ is a 2ϵ-net in C. Therefore every element of C is the limit of a sequence from M^n and hence C is contained in M^n. □

10.5 Corollary. *Let $\mathcal{M} \preceq \mathcal{N}$ be L-structures and let A be a subset of M. Then*

$$\mathrm{dcl}_{\mathcal{M}}(A) = \mathrm{dcl}_{\mathcal{N}}(A) \subseteq \mathrm{acl}_{\mathcal{N}}(A) = \mathrm{acl}_{\mathcal{M}}(A).$$

Proof Proposition 10.4 shows that if $C \subseteq N^n$ is compact and is definable in $\mathcal{N}$ over A, then $C \subseteq M^n$. Obviously C is then definable in $\mathcal{M}$ over A. It follows that $\mathrm{dcl}_{\mathcal{N}}(A) \subseteq \mathrm{dcl}_{\mathcal{M}}(A)$ and $\mathrm{acl}_{\mathcal{N}}(A) \subseteq \mathrm{acl}_{\mathcal{M}}(A)$.

For the opposite containment, suppose $C \subseteq M^n$ is compact and definable in $\mathcal{M}$ over A. We want to show that C is definable in $\mathcal{N}$ over A. Let $P(x) = \mathrm{dist}(x, C)$ for all $x \in M^n$. Since P is definable in $\mathcal{M}$ over A, by Proposition 9.8 there is a predicate $Q \colon N^n \to [0,1]$ that is definable in $\mathcal{N}$ over A and that extends P. Moreover we have $(\mathcal{M}, P) \preceq (\mathcal{N}, Q)$. Let $D \subseteq N^n$ be the zero set of Q. By Theorem 9.12 we have that $Q(x) = \mathrm{dist}(x, D)$ for all $x \in N^n$. It follows from Proposition 10.4 that $D = C$. Therefore C is definable in $\mathcal{N}$ over A. It follows that $\mathrm{dcl}_{\mathcal{M}}(A) \subseteq \mathrm{dcl}_{\mathcal{N}}(A)$ and $\mathrm{acl}_{\mathcal{M}}(A) \subseteq \mathrm{acl}_{\mathcal{N}}(A)$. □

An alternative way to approach *definable closure* and *algebraic closure* in metric structures would be to consider compact *zerosets* rather than definable sets. The next result shows that the concepts would be the same, as long as one takes care to work in saturated models.

10.6 Proposition. *Let $\mathcal{M}$ be an ω_1-saturated L-structure and A a subset of M. For compact subsets C of M^n the following are equivalent:*
(1) C is a zeroset in $\mathcal{M}$ over A.
(2) C is definable in $\mathcal{M}$ over A.

Proof Obviously (2) implies (1). For the converse, let $P \colon M^n \to [0,1]$ be a predicate definable in $\mathcal{M}$ over A whose zero set is C.

Given $\epsilon > 0$, let $F \subseteq C$ be a finite $\epsilon/2$-net in C. We claim there exists $\delta > 0$ such that any a satisfying $P(a) \leq \delta$ must lie within ϵ of some element of F. Otherwise we may use the ω_1-saturation of $\mathcal{M}$ to obtain an element a of M^n such that $P(a) \leq 1/k$ for every $k \geq 1$ while $d(a,c) \geq \epsilon$ for all $c \in F$, which is impossible.

The existence of such a $\delta > 0$ for each $\epsilon > 0$ shows that P verifies (2) in Proposition 9.19. Hence C is definable in $\mathcal{M}$ over A. □

The following result gives alternative characterizations of definability.

10.7 Exercise. *Let $\mathcal{M}$ be an L-structure, $A \subseteq M$, and $a \in M^n$. Statements (1),(2) and (3) are equivalent:*
(1) a is definable in $\mathcal{M}$ over A.
(2) For any $\mathcal{N} \succeq \mathcal{M}$ the only realization of $\mathrm{tp}_{\mathcal{M}}(a/A)$ in $\mathcal{N}$ is a.
(3) For any $\epsilon > 0$ there is an $L(A)$-formula $\varphi(x)$ and $\delta > 0$ such that $\varphi^{\mathcal{M}}(a) = 0$ and the diameter of $\{b \in M^n \mid \varphi^{\mathcal{M}}(b) < \delta\}$ is $\leq \epsilon$.
If $\mathcal{N}$ is any fixed ω_1-saturated elementary extension of $\mathcal{M}$, then (1) is equivalent to:
(4) The only realization of $\mathrm{tp}_{\mathcal{M}}(a/A)$ in $\mathcal{N}$ is a.

The next result gives alternative characterizations of algebraicity:

10.8 Exercise. *Let $\mathcal{M}$ be an L-structure, $A \subseteq M$, and $a \in M^n$. Statements (1),(2) and (3) are equivalent:*
(1) a is algebraic in $\mathcal{M}$ over A.
(2) For any $\mathcal{N} \succeq \mathcal{M}$, every realization of $\mathrm{tp}_{\mathcal{M}}(a/A)$ in $\mathcal{N}$ is in M^n.
(3) For any $\epsilon > 0$ there is an $L(A)$-formula $\varphi(x)$ and $\delta > 0$ such that $\varphi^{\mathcal{M}}(a) = 0$ and the set $\{b \in M^n \mid \varphi^{\mathcal{M}}(b) < \delta\}$ has a finite ϵ-net.
(4) For any $\mathcal{N} \succeq \mathcal{M}$, the set of realizations of $\mathrm{tp}_{\mathcal{M}}(a/A)$ in $\mathcal{N}$ is compact.
If $\mathcal{N}$ is any fixed ω_1-saturated elementary extension of $\mathcal{M}$, then (1) is equivalent to:
(5) The set of realizations of $\mathrm{tp}_{\mathcal{M}}(a/A)$ in $\mathcal{N}$ is compact.
If $\mathcal{N}$ is any fixed κ-saturated elementary extension of $\mathcal{M}$, with κ uncountable, then (1) is equivalent to:
(6) The set of realizations of $\mathrm{tp}_{\mathcal{M}}(a/A)$ in $\mathcal{N}$ has density character $< \kappa$.

10.9 Remark. The previous result shows that $\mathrm{acl}_{\mathcal{M}}(A)$ is the same as the *bounded* closure of A in $\mathcal{M}$. (Equivalence of (1) and (6) in highly saturated models.)

The usual basic properties of algebraic and definable closure in first-order logic hold in this more general setting. There is one (unsurprising) modification needed: if a is algebraic (resp., definable) over A in

$\mathcal{M}$, there is a *countable* subset A_0 of A such that a is algebraic (resp., definable) over A_0 in $\mathcal{M}$. However, it may be impossible to satisfy this property with A_0 finite. (For example, let A consist of the elements of a Cauchy sequence whose limit is a. Then a is definable over A but it is not necessarily even algebraic over a finite subset of A.) There is also one new feature: the algebraic (resp., definable) closure of a set is the same as the algebraic closure of any dense subset.

In what follows we fix an L-structure $\mathcal{M}$ and A, B are subsets of M. We write acl instead of $\mathrm{acl}_{\mathcal{M}}$.

10.10 Exercise. *(Properties of* dcl*)*
(1) $A \subseteq \mathrm{dcl}(A)$*;*
(2) if $A \subseteq \mathrm{dcl}(B)$ *then* $\mathrm{dcl}(A) \subseteq \mathrm{dcl}(B)$*;*
(3) if $a \in \mathrm{dcl}(A)$ *then there exists a countable set* $A_0 \subseteq A$ *such that* $a \in \mathrm{dcl}(A_0)$*;*
(4) if A *is a dense subset of* B*, then* $\mathrm{dcl}(A) = \mathrm{dcl}(B)$.

10.11 Proposition. *(Properties of* acl*)*
(1) $A \subseteq \mathrm{acl}(A)$*;*
(2) if $A \subseteq \mathrm{acl}(B)$ *then* $\mathrm{acl}(A) \subseteq \mathrm{acl}(B)$*;*
(3) if $a \in \mathrm{acl}(A)$ *then there exists a countable set* $A_0 \subseteq A$ *such that* $a \in \mathrm{acl}(A_0)$*;*
(4) if A *is a dense subset of* B*, then* $\mathrm{acl}(A) = \mathrm{acl}(B)$.

10.12 Proposition. *If* $\alpha \colon A \to B$ *is an elementary map, then it extends to an elementary map* $\alpha' \colon \mathrm{acl}(A) \to \mathrm{acl}(B)$*. Moreover, if* α *is surjective, then so is* α'*.*

Proofs of the last two results are given in [26].

11 Imaginaries

In this section we explain how to add *finitary imaginaries* to a metric structure. This construction is the first stage of a natural generalization of the $\mathcal{M}^{eq}$ construction in ordinary first-order model theory.

There are several different ways to look at this construction. We first take a "geometric" point of view, based on forming the quotient of a pseudometric space. This is a generalization to the metric setting of taking the quotient of a set X by an equivalence relation E and connecting X to X/E by the quotient map π_E which takes each element of X to its E-class.

Let L be any signature for metric structures and let $\rho(x,y)$ be an L-formula in which x and y are n-tuples of variables. If $\mathcal{M}$ is any L-structure, we think of $\rho^{\mathcal{M}}$ as being a function of two variables on M^n, and we are interested in the case where this function is a pseudometric on M^n. (Note: to consider a pseudometric that is definable in $\mathcal{M}$ over $\emptyset$, but is not itself the interpretation of a formula, one should first pass to an extension by definitions in which it is the interpretation of a predicate.)

We define a signature L_ρ that extends L by adding a new sort, which we denote by M' with metric ρ', and a new n-ary function symbol π_ρ, which is to be interpreted by functions from M^n into M'. The modulus of uniform continuity specified by L_ρ for π_ρ is Δ_ρ. (Recall that Δ_ρ is a modulus of uniform continuity for $\rho^{\mathcal{M}}$ in every L-structure $\mathcal{M}$. See Theorem 3.5.)

Now suppose $\mathcal{M}$ is any L-structure in which $\rho^{\mathcal{M}}$ is a pseudometric on M^n. We expand $\mathcal{M}$ to an L_ρ-structure by interpreting (M',ρ') to be the completion of the quotient metric space of $(M^n,\rho^{\mathcal{M}})$ and by interpreting π_ρ to be the canonical quotient mapping from M^n into M'. This expansion of $\mathcal{M}$ will be denoted by $\mathcal{M}_\rho$. Saying that $\mathcal{M}_\rho$ is an L_ρ-structure requires checking that Δ_ρ is a modulus of uniform continuity for the interpretation of π_ρ. We indicate why this is true: suppose $\epsilon > 0$ and $x, x' \in M^n$ satisfy $d(x,x') < \Delta_\rho(\epsilon)$. Then

$$\rho'(\pi_\rho(x),\pi_\rho(x')) = \rho(x,x') = |\rho(x,x') - \rho(x',x')| \leq \epsilon$$

as desired.

Let T_ρ be the L_ρ-theory consisting of the following conditions:

(1) $\sup_x \rho(x,x) = 0$;
(2) $\sup_x \sup_y |\rho(x,y) - \rho(y,x)| = 0$;
(3) $\sup_x \sup_y \sup_{y'} \big(\rho(x,y) \dot- \min(\rho(x,y') + \rho(y',y), 1)\big) = 0$;
(4) $\sup_x \sup_y |\rho'(\pi_\rho(x),\pi_\rho(y)) - \rho(x,y)| = 0$;
(5) $\sup_z \inf_x \rho'(z,\pi_\rho(x)) = 0$.

11.1 Theorem. *(1) For every L-structure $\mathcal{M}$ in which $\rho^{\mathcal{M}}$ is a pseudometric, the expansion $\mathcal{M}_\rho$ is a model of T_ρ.*
(2) If $\mathcal{N}$ is any model of T_ρ, with $\mathcal{M}$ its reduct to L, then $\rho^{\mathcal{M}}$ is a pseudometric and $\mathcal{N}$ is isomorphic to $\mathcal{M}_\rho$ by an isomorphism that is the identity on $\mathcal{M}$.

Proof (1) is obvious. If $\mathcal{N}$ is any model of T_ρ, with $\mathcal{M}$ its reduct to L, the first three conditions in T_ρ ensure that $\rho^{\mathcal{M}}$ is a pseudometric.

Statement (4) ensures that the interpretation of π_ρ together with the metric ρ' on its range is isomorphic to the canonical quotient of the pseudometric space $(M^n, \rho^{\mathcal{M}})$ by an isomorphism that is the identity on M^n. Statement (5) ensures that the range of the interpretation of π_ρ is dense in (M', ρ'). □

To keep notation simple we have limited our discussion to the situation where L is a 1-sorted signature. If it is many-sorted, we may replace M^n by any finite cartesian product of sorts and carry out exactly the same construction. In particular, the process of adding imaginary sorts can be iterated.

Canonical parameters for formulas

Now we use the quotient construction above to add canonical parameters for any L-formula $\varphi(u, x)$, where u is an m-tuple of variables and x is an n-tuple of variables (thought of as the parameters). We take $\rho(x, y)$ to be the L-formula $\sup_u |\varphi(u, x) - \varphi(u, y)|$. Note that for any L-structure $\mathcal{M}$ we have that $\rho^{\mathcal{M}}$ is a pseudometric on M^n. Thus we may carry out the construction above, obtaining the signature L_ρ and the expansion $\mathcal{M}_\rho$ of any L-structure $\mathcal{M}$ to a uniquely determined model of T_ρ. Let $\widehat{\varphi}(u, z)$ be the L_ρ-formula

$$\inf_y \big(\varphi(u, y) + \rho'(z, \pi_\rho(y))\big).$$

11.2 Proposition. *If $\mathcal{N}$ is any model of T_ρ, then*

$$\mathcal{N} \models \sup_u \sup_x |\widehat{\varphi}(u, \pi_\rho(x)) - \varphi(u, x)| = 0 \text{ and}$$

$$\mathcal{N} \models \sup_w \sup_z \big|\rho'(w, z) - \sup_u |\widehat{\varphi}(u, w) - \widehat{\varphi}(u, z)|\big| = 0.$$

Proof Let $\mathcal{N}$ be any model of T_ρ and let $\mathcal{M}$ be its reduct to L. Let (M', ρ') be the sort of $\mathcal{N}$ that is added when expanding $\mathcal{M}$ to an L_ρ-structure.

To prove the first statement, take any $x, y \in M^n$. Then for any $u \in M^m$ we have

$$\begin{aligned} \varphi^{\mathcal{M}}(u, y) + \rho'(\pi_\rho^{\mathcal{N}}(x), \pi_\rho^{\mathcal{N}}(y)) &= \varphi^{\mathcal{M}}(u, y) + \rho^{\mathcal{M}}(x, y) \\ &\geq \varphi^{\mathcal{M}}(u, y) + |\varphi^{\mathcal{M}}(u, x) - \varphi^{\mathcal{M}}(u, y)| \\ &\geq \varphi^{\mathcal{M}}(u, x). \end{aligned}$$

Taking the inf over $y \in M^n$ shows that $\widehat{\varphi}^{\mathcal{N}}(u, \pi_\rho^{\mathcal{N}}(x)) \geq \varphi^{\mathcal{M}}(u,x)$. On the other hand, taking $y = x$ in the definition of $\widehat{\varphi}^{\mathcal{N}}$ shows that $\widehat{\varphi}^{\mathcal{N}}(u, \pi_\rho^{\mathcal{N}}(x)) \leq \varphi^{\mathcal{M}}(u,x)$.

For the second statement, note that it suffices to consider w, z in the range of $\pi_\rho^{\mathcal{N}}$, since it is dense in M'. When $w = \pi_\rho^{\mathcal{N}}(x)$ and $z = \pi_\rho^{\mathcal{N}}(y)$, the equality to be proved follows easily from the first statement in the Proposition and the definition of $\rho(x,y)$. □

11.3 Remark. Let $\mathcal{M}$ be any L-structure and $\mathcal{N} = \mathcal{M}_\rho$ its canonical expansion to a model of T_ρ. As before, denote the extra sort of this expansion by (M', ρ'). The results above show that M' is a space of canonical parameters for the predicate $\varphi^{\mathcal{M}}(u,x)$ and that $\widehat{\varphi}^{\mathcal{N}}$ is the predicate resulting from $\varphi^{\mathcal{M}}$ by the identification of each $x \in M^n$ with its associated canonical parameter $z = \pi_\rho^{\mathcal{N}}(x)$.

11.4 Remark. Note that in any model $\mathcal{N}$ of T_ρ the function $\pi_\rho^{\mathcal{N}}$ is definable from $\widehat{\varphi}^{\mathcal{N}}$ and $\varphi^{\mathcal{N}}$ in the sense of Definition 9.22. Indeed, when $\mathcal{N} \models T_\rho$ we have

$$\mathcal{N} \models \sup_x \sup_z \big|\rho'(\pi_\rho(x), z) - \sup_u |\varphi(u,x) - \widehat{\varphi}(u,z)|\big| = 0.$$

(Proof: Specialize w to $\pi_\rho(x)$ in the second statement in Proposition 11.2 and then use the first statement to replace $\widehat{\varphi}(u, \pi_\rho(x))$ by $\varphi(u,x)$.)

In the metric setting, the construction of $\mathcal{M}^{eq}$ should allow expansions corresponding to extensions by definitions as well as those corresponding to quotients by definable (over $\emptyset$) pseudometrics.

In fact, the full construction of $\mathcal{M}^{eq}$ requires the addition of more sorts than are described here. In particular, for some uses in stability theory one needs to add imaginaries that provide canonical parameters for definable predicates that depend on *countably many* parameters. (What we describe here only covers the case where the predicate depends on *finitely many* parameters.) See [6, end of Section 5] for a sketch of how this is done.

12 Omitting types and ω-categoricity

In this section we assume that the signature L has only a countable number of nonlogical symbols.

Let T be a complete L-theory. We emphasize here our point of view that models of T are always complete for their underlying metric(s). That is especially significant for the meaning of properties such as categoricity and omitting types.

12.1 Definition. Let κ be a cardinal $\geq \omega$. We say T is κ-categorical if whenever $\mathcal{M}$ and $\mathcal{N}$ are models of T having density character equal to κ, one has that $\mathcal{M}$ is isomorphic to $\mathcal{N}$.

One of the main goals of this section is to state a characterization of ω-categoricity for complete theories in continuous logic, extending the Ryll-Nardzewski Theorem from first-order model theory. We note that Ben Yaacov has proved the analogue for this setting of Morley's Theorem concerning uncountable categoricity. (See [4].)

Let p be an n-type over $\emptyset$ for T; that is, $p \in S_n(T)$. If $\mathcal{M}$ is a model of T, we let $p(\mathcal{M})$ denote the set of realizations of p in $\mathcal{M}$.

12.2 Definition. Let $p \in S_n(T)$. We say that p is *principal* if for every model $\mathcal{M}$ of T, the set $p(\mathcal{M})$ is definable in $\mathcal{M}$ over $\emptyset$.

The following Lemma shows that for $p \in S_n(T)$ to be principal, it suffices for there to exist *some* model of T in which the set of realizations of p is a nonempty definable set. It also shows that a principal type is realized in every model of T.

12.3 Lemma. *Let $p \in S_n(T)$. Suppose there exists a model $\mathcal{M}$ of T such that $p(\mathcal{M})$ is nonempty and definable in $\mathcal{M}$ over $\emptyset$. Then for any model $\mathcal{N}$ of T, the set $p(\mathcal{N})$ is nonempty and definable in $\mathcal{N}$ over $\emptyset$.*

Proof Since T is complete, any two models of T have isomorphic elementary extensions. Therefore it suffices to consider the case in which one of the structures is an elementary extension of the other.

First suppose $\mathcal{M} \preceq \mathcal{N}$ and that $p(\mathcal{N})$ is nonempty and definable in $\mathcal{N}$ over $\emptyset$. Since $p(\mathcal{M}) = p(\mathcal{N}) \cap M^n$, Theorem 9.12 and Proposition 9.18 yield that $p(\mathcal{M})$ is nonempty and definable in $\mathcal{M}$ over $\emptyset$.

Now suppose $\mathcal{M} \preceq \mathcal{N}$ and that $p(\mathcal{M})$ is nonempty and definable in $\mathcal{M}$ over $\emptyset$. Let $P(x) = \operatorname{dist}(x, p(\mathcal{M}))$, so P is a definable predicate in $\mathcal{M}$ (over $\emptyset$). Use Proposition 9.8 to obtain a predicate Q on N^n so that $(\mathcal{M}, P) \preceq (\mathcal{N}, Q)$. By Theorem 9.12 the predicate Q satisfies $Q(x) = \operatorname{dist}(x, D)$ for all $x \in N^n$, where D is the zeroset of Q. Therefore, it suffices to show that $p(\mathcal{N})$ is the zeroset of Q.

First we prove that the zeroset of Q is contained in $p(\mathcal{N})$. Suppose

$\varphi(x)$ is any L-formula for which the condition $\varphi = 0$ is in p. Fix $\epsilon > 0$; by Theorem 3.5 there exists $\delta > 0$ such that whenever $x, y \in M^n$ satisfy $d(x, y) < \delta$, then $|\varphi^{\mathcal{M}}(x) - \varphi^{\mathcal{M}}(y)| \leq \epsilon$. We may assume $\delta < 1$. If $x \in M^n$ satisfies $\mathrm{dist}(x, p(\mathcal{M})) < \delta$, there must exist $y \in M^n$ realizing p with $d(x, y) < \delta$ and hence

$$\varphi^{\mathcal{M}}(x) = |\varphi^{\mathcal{M}}(x) - \varphi^{\mathcal{M}}(y)| \leq \epsilon.$$

That is,

$$(\mathcal{M}, P) \models \sup_x (\min(\delta \dot{-} P(x), \varphi(x) \dot{-} \epsilon)) = 0.$$

It follows that

$$(\mathcal{N}, Q) \models \sup_x (\min(\delta \dot{-} Q(x), \varphi(x) \dot{-} \epsilon)) = 0.$$

Since ϵ was arbitrary and $\varphi = 0$ was an arbitrary condition in p, this shows that the zeroset of Q is contained in $p(\mathcal{N})$.

Finally, we need to show that Q is identically zero on $p(\mathcal{N})$. By construction, there is a sequence $(\varphi_n(x))$ of L-formulas such that

$$|Q(x) - \varphi_n^{\mathcal{N}}(x)| \leq 1/n$$

for all $x \in N^n$ and all $n \geq 1$. Since $p(\mathcal{M})$ is nonempty, we may take $x \in M^n$ realizing p; since $Q(x) = P(x) = \mathrm{dist}(x, p(\mathcal{M})) = 0$, we see that the condition $\varphi_n(x) \leq 1/n$ is in $p(x)$ for all $n \geq 1$. For any $x \in p(\mathcal{N})$ we therefore have

$$Q(x) \leq |Q(x) - \varphi_n^{\mathcal{N}}(x)| + \varphi_n^{\mathcal{N}}(x) \leq 2/n$$

for all $n \geq 1$. Therefore $Q(x) = 0$ for any x in $p(\mathcal{M})$. □

12.4 Proposition. *Let $p \in S_n(T)$. Then p is principal if and only if the logic topology and the d-metric topology agree at p.*

Proof Let $\mathcal{M}$ be an ω-saturated model of T. Since p is realized in $\mathcal{M}$, Lemma 12.3 implies that p is principal if and only if $p(\mathcal{M})$ is definable in $\mathcal{M}$ over $\emptyset$. We apply Proposition 9.19, taking $D = p(\mathcal{M})$ and $A = \emptyset$. This yields that p is principal if and only if for each $m \geq 1$ there is an L-formula $\varphi_m(x)$ and $\delta_m > 0$ such that $\varphi_m = 0$ is in p and any $q \in S_n(T)$ that contains the condition $\varphi_m \leq \delta_m$ must satisfy $d(q, p) \leq 1/m$.

So, when p is principal and $m \geq 1$, we have a logic open neighborhood of p, namely $[\varphi_m < \delta_m]$, that is contained in the $1/m$-ball around p.

On the other hand, suppose $[\psi < \delta]$ is a basic logic neighborhood of

p that is contained in the $1/m$-ball around p. There exists $0 < \eta < \delta$ such that the condition $\psi \leq \eta$ is in p. Taking φ_m to be the formula $\psi \dot{-} \eta$ and δ_m to satisfy $0 < \delta_m < \delta - \eta$, we have that $\varphi_m = 0$ is in p and that any $q \in S_n(T)$ that contains the condition $\varphi_m \leq \delta_m$ must satisfy $d(q,p) \leq 1/m$. If this is possible for every $m \geq 1$, then p must be principal. □

The fact that the metric is included as a predicate in L allows us to characterize principal types by a formally weaker topological property than the one in the previous Proposition:

12.5 Proposition. *Let $p \in S_n(T)$. Then p is principal if and only if the ball $\{q \in S_n(T) \mid d(q,p) \leq \epsilon\}$ has nonempty interior in the logic topology, for each $\epsilon > 0$.*

Proof (⇒) This follows from Proposition 12.4.

(⇐) Suppose $[\psi < \delta]$ is a nonempty basic open set contained in the ϵ-ball around $p \in S_n(T)$. We may assume $\delta \leq \epsilon$, since $[\frac{1}{k}\psi < \frac{1}{k}\delta] = [\psi < \delta]$ for all k. Choose η such that $0 < \eta < \delta$ and $[\psi \leq \eta]$ is nonempty. Consider the formula

$$\varphi(x) = \inf_y \max(\psi(y) \dot{-} \eta, d(x,y) \dot{-} \epsilon).$$

The condition $\varphi = 0$ is in p. This is because $[\psi \leq \eta]$ is nonempty and is contained in the ϵ-ball around p. Furthermore, the basic open set $[\varphi < \delta - \eta]$ is contained in the $(2\epsilon + \delta)$-ball around p. Taking ϵ arbitrarily small gives the desired result, by Proposition 12.4. □

12.6 Theorem. *(Omitting Types Theorem, local version) Let T be a complete theory in a countable signature, and let $p \in S_n(T)$. The following statements are equivalent:*

(1) *p is principal.*
(2) *p is realized in every model of T.*

Proof (1) ⇒ (2). Since p must be realized in *some* model of T, Lemma 12.3 shows that a principal type is realized in *all* models of T.

(2) ⇒ (1). We sketch a proof of the contrapositive. Suppose p is not principal. By Proposition 12.5, there exists $\epsilon > 0$ such that the ϵ-ball $\{q \in S_n(T) \mid d(q,p) \leq \epsilon\}$ has empty interior in the logic topology. That is, for any L-formula $\varphi(x)$ and any $\delta > 0$, the logic neighborhood $[\varphi < \delta]$ is either empty or contains a type q such that $d(q,p) > \epsilon$. An argument

as in the usual proof of the omitting types theorem in classical first-order model theory yields a countable L-prestructure $\mathcal{M}_0$ satisfying the theory T^+, such that any n-type q realized in $\mathcal{M}_0$ satisfies $d(q,p) > \epsilon$. Let $\mathcal{M}$ be the completion of $\mathcal{M}_0$. It follows that any type q realized in $\mathcal{M}$ satisfies $d(q,p) \geq \epsilon$, and hence $\mathcal{M}$ is a model of T in which p is not realized. □

12.7 Definition. A model $\mathcal{M}$ of T is *atomic* if every n-type realized in $\mathcal{M}$ is principal.

12.8 Proposition. *Let $\mathcal{M}$ be a model of T and let $p(x_1,\dots,x_m) = \mathrm{tp}_{\mathcal{M}}(a_1,\dots,a_m)$ for a sequence $a_1,\dots,a_m$ in M. Let $n > m$ and suppose $q(x_1,\dots,x_n) \in S_n(T)$ extends p. If q is principal, then for each $\epsilon > 0$ there exist $(b_1,\dots,b_n)$ realizing q in M and satisfying $d(a_j,b_j) \leq \epsilon$ for all $j = 1,\dots,m$.*

Proof Let $D \subseteq M^m$ and $E \subseteq M^n$ be the sets of realizations of p and q respectively in $\mathcal{M}$. Because q is principal, both D and E are definable in $\mathcal{M}$ over $\emptyset$. Hence the function F defined for $x \in M^m$ by

$$F(x) = \inf_y |\mathrm{dist}(x,D) - \mathrm{dist}((x,y),E)|$$

is a predicate definable in $\mathcal{M}$ over $\emptyset$. Here y ranges over M^{n-m}. If $\mathcal{N}$ is an ω_1-saturated elementary extension of $\mathcal{M}$, the sequence $(a_1,\dots,a_m)$ can be extended in $\mathcal{N}$ to a realization of q. If we extend F to N^m using the same definition, this shows that $F(a_1,\dots,a_m) = 0$ in $\mathcal{N}$. By Lemma 9.6 we have that $F(a_1,\dots,a_m) = 0$ in $\mathcal{M}$. So there exist $c \in M^{n-m}$ such that the n-tuple $(a_1,\dots,a_m,c)$ has distance $\leq \epsilon$ to E in M^n. This completes the proof. □

12.9 Corollary. *Let $\mathcal{M}$ be a separable atomic model of T and let $\mathcal{N}$ be any other model of T. Let $(a_1,\dots,a_m)$ realize the same type in $\mathcal{M}$ that $(b_1,\dots,b_m)$ realizes in $\mathcal{N}$. Then, for each $\epsilon > 0$ there exists an elementary embedding F from $\mathcal{M}$ into $\mathcal{N}$ such that $d(b_j, F(a_j)) \leq \epsilon$ for all $j = 1,\dots,m$. Furthermore, if $\mathcal{N}$ is also separable and atomic, then F can be taken to be an isomorphism from $\mathcal{M}$ onto $\mathcal{N}$.*

Proof Extend $a_1,\dots,a_m$ to an infinite sequence (a_k) that is dense in M. Let (δ_k) be a sequence of positive real numbers whose sum is less than ϵ. By induction on $n \geq 0$ we use the previous Proposition to generate sequences c_n in N^{m+n} with the following properties: (1) $c_0 =$

$(b_1, \ldots, b_m)$; (2) c_n realizes the same type in $\mathcal{N}$ that $(a_1, \ldots, a_{m+n})$ realizes in $\mathcal{M}$; (3) the first $m+n$ coordinates of c_{n+1} are at a distance less than δ_n away from the corresponding coordinates of c_n.

It follows that for each j, the sequence of j^{th} coordinates of c_n is a Cauchy sequence in N. Let its limit be d_j. Continuity of formulas ensures that the map taking a_j to d_j for each j is an elementary map. Therefore it extends to the desired elementary embedding of $\mathcal{M}$ into $\mathcal{N}$.

A "back-and-forth" version of the same argument proves the final statement in the Corollary. □

The following result is the analogue of the Ryll-Nardzewski Theorem in this setting:

12.10 Theorem. *Let T be a complete theory in a countable signature. The following statements are equivalent:*

(1) *T is ω-categorical;*
(2) *For each $n \geq 1$, every type in $S_n(T)$ is principal;*
(3) *For each $n \geq 1$, the metric space $(S_n(T), d)$ is compact.*

Proof (1) ⇒ (2) This is immediate from the Omitting Types Theorem.

(2) ⇒ (1) Condition (2) implies that every model of T is atomic. Therefore, Corollary 12.9 (especially, the last sentence) yields that any two separable models of T are isomorphic.

(2) ⇔ (3) By Proposition 12.4, statement (2) is equivalent to saying that the logic topology and the d-metric topology are identical on $S_n(T)$ for every n. Since the logic topology is compact, this shows that (2) implies (3). On the other hand, if the metric space $(S_n(T), d)$ is compact, then its topology must agree with the logic topology, since both topologies are compact and Hausdorff, and one is coarser than the other. □

12.11 Corollary. *Suppose T is ω-categorical and $\mathcal{M}$ is the separable model of T. Then $\mathcal{M}$ is* strongly ω-near-homogeneous *in the following sense: if $a, b \in M^n$, then for every $\epsilon > 0$ there is an automorphism F of $\mathcal{M}$ such that*

$$d(F(a), b) \leq d(\mathrm{tp}(a), \mathrm{tp}(b)) + \epsilon.$$

Proof By Theorem 12.10 we see that $\mathcal{M}$ is atomic. Let

$$r = d(\mathrm{tp}(a), \mathrm{tp}(b)).$$

First consider the case where $r = 0$. Corollary 12.9 yields the existence of an automorphism F of $\mathcal{M}$ with the desired properties. Now suppose $r > 0$. Since $\mathcal{M}$ is the unique separable model of T, there exist $a', b' \in M^n$ such that $\mathrm{tp}(a') = \mathrm{tp}(a)$, $\mathrm{tp}(b') = \mathrm{tp}(b)$, and $d(a', b') = r$. Applying the $r = 0$ case to a, a' and to b, b', for each $\epsilon > 0$ we get automorphisms F_a, F_b of $\mathcal{M}$ such that $d(F_a(a), a') < \epsilon/2$ and $d(F_b(b), b') < \epsilon/2$. Taking $F = F_b^{-1} \circ F_a$ gives an automorphism with

$$\begin{aligned} d(F(a), b) &= d(F_a(a), F_b(b)) \\ &\leq d(F_a(a), a') + d(a', b') + d(F_b(b), b') \\ &< r + \epsilon \end{aligned}$$

as desired. □

12.12 Remark. In the previous result, note that for any automorphism F of $\mathcal{M}$, $d(F(a), b) \geq d(\mathrm{tp}(a), \mathrm{tp}(b))$ since $F(a)$ and a realize the same type. Moreover, Example 17.7 shows that in the setting of the previous result, we need not be able to find an automorphism F of $\mathcal{M}$ that satisfies $d(F(a), b) = d(\mathrm{tp}(a), \mathrm{tp}(b))$; so this result gives the strongest possible kind of homogeneity for ω-categorical metric structures, in general. (Example 17.7 notes that the theory of the Banach lattice L^p is ω-categorical, but that there exist elements f, g realizing the same type but are such that there is no automorphism of L^p taking f to g.)

12.13 Corollary. *Suppose $L \subseteq L'$ are countable signatures, T' is a complete theory in L' and T is its restriction to L. If T' is ω-categorical, then so is T.*

Proof The restriction map (discarding formulas not in L) defines a map from $S_n(T')$ onto $S_n(T)$ that is contractive with respect to the d-metrics. Hence it preserves compactness. □

12.14 Remark. Example 17.7 shows that ω-categoricity is *not* necessarily preserved under the addition of designated elements to the language, in contrast to what happens in classical first-order model theory. There can exist pairs (a, b) realizing a principal type in some model, but such that $\mathrm{tp}(b/a)$ is not principal. This indicates a complication in the model theory of metric structures that is not completely understood, and that affects a number of important aspects of the theory (including, for example, superstability).

13 Quantifier elimination

Fix a signature L and an L-theory T. We give basic definitions and state (but do not prove) some results around quantifier elimination in continuous logic. The proofs are similar to those in [24, pages 84–91].

13.1 Definition. An L-formula $\varphi(x_1, \dots, x_n)$ is *approximable in T by quantifier-free formulas* if for every $\epsilon > 0$ there is a quantifier-free L-formula $\psi(x_1, \dots, x_n)$ such that for all $\mathcal{M} \models T$ and all $a_1, \dots, a_n \in M$, one has

$$|\varphi^{\mathcal{M}}(a_1, \dots, a_n) - \psi^{\mathcal{M}}(a_1, \dots, a_n)| \leq \epsilon.$$

13.2 Proposition. *Let $\varphi(x_1, \dots, x_n)$ be an L-formula. The following statements are equivalent.*

(1) *φ is approximable in T by quantifier-free formulas;*
(2) *Whenever we are given*

- *models $\mathcal{M}$ and $\mathcal{N}$ of T;*
- *substructures $\mathcal{M}_0 \subseteq \mathcal{M}$ and $\mathcal{N}_0 \subseteq \mathcal{N}$;*
- *an isomorphism Φ from $\mathcal{M}_0$ onto $\mathcal{N}_0$; and*
- *elements $a_1, \dots, a_n$ of $\mathcal{M}_0$;*

we have

$$\varphi^{\mathcal{M}}(a_1, \dots, a_n) = \varphi^{\mathcal{N}}(\Phi(a_1), \dots, \Phi(a_n)).$$

Moreover, for the implication (2)⇒(1) *it suffices to assume* (2) *only for the cases in which $\mathcal{M}_0$ and $\mathcal{N}_0$ are finitely generated.*

13.3 Definition. An L-theory T *admits quantifier elimination* if every L-formula is approximable in T by quantifier-free formulas.

13.4 Remark. (1) Let T be an L-theory, and let $L(C)$ be an extension of L by constants. If T admits quantifier elimination in L, then T admits quantifier elimination in $L(C)$.
(2) Let $T \subseteq T'$ be theories in a signature L. If T admits quantifier elimination in L, then T' admits quantifier elimination in L.

13.5 Lemma. *Suppose that T is an L-theory and that every restricted L-formula of the form $\inf_x \varphi$, with φ quantifier-free, is approximable in T by quantifier-free formulas. Then T admits quantifier elimination.*

13.6 Proposition. *Let T be an L-theory. Then the following statements are equivalent:*

(1) *T admits quantifier elimination;*
(2) *If $\mathcal{M}$ and $\mathcal{N}$ are models of T, then every embedding of a substructure of $\mathcal{M}$ into $\mathcal{N}$ can be extended to an embedding of $\mathcal{M}$ into an elementary extension of $\mathcal{N}$.*

Moreover, if $\operatorname{card}(L) \leq \kappa$*, then in statement (2) it suffices to consider models $\mathcal{M}$ of density character $\leq \kappa$.*

14 Stability and independence

In this section we sketch three general approaches to stability in metric structures. The first one is based on measuring the size of the type spaces $S_1(T_A)$ in various ways. The second one comes from properties of the notion of independence obtained from non-dividing. The third one comes from definability of types. In all three cases we give clear statements of the definitions and the basic results (which we need in later sections), but we only give some of the proofs. In spite of having these different approaches to stability, there is, in the end, only one notion of stability, as we discuss.

Throughout this section T is a complete L-theory, κ is a cardinal $> \operatorname{card}(L)$, and λ is an infinite cardinal. As is usual, we often denote $\operatorname{card}(L)$ also by $|T|$; recall that this is the least *infinite* cardinal $\geq$ the number of nonlogical symbols in L. When $\mathcal{M}$ is a model of T and $A \subseteq M$, recall that T_A is the theory of $(\mathcal{M}, a)_{a \in A}$. We take x and y to be finite sequences of distinct variables; usually $x = x_1, \ldots, x_n$.

We begin with an approach to stability based simply on the cardinality of the type spaces $S_1(T_A)$.

14.1 Definition. We say that T is λ-*stable with respect to the discrete metric* if for any $\mathcal{M} \models T$ and any $A \subseteq M$ of cardinality $\leq \lambda$, the set $S_1(T_A)$ has cardinality $\leq \lambda$. We say that T is *stable with respect to the discrete metric* if T is λ-stable with respect to the discrete metric for some λ.

14.2 Definition. Let $\mathcal{U}$ be a κ-universal domain for T. Let $\varphi(x, y)$, $\psi(x, y)$ be formulas such that the conditions $\varphi(x, y) = 0$ and $\psi(x, y) = 0$ are contradictory in $\mathcal{U}$. Since $\mathcal{U}$ is κ-saturated, there exists some $\epsilon > 0$ such that $\{\varphi(x, y) = 0, \psi(x', y) = 0, d(x, x') \leq \epsilon\}$ is not satisfiable in $\mathcal{U}$.

If $p(x)$ is any satisfiable partial type over a small subset of $\mathcal{U}$, we define the rank $R(p, \varphi, \psi, 2)$ inductively, in the usual manner. First we define a relation $R(p, \varphi, \psi, 2) \geq \alpha$ by induction on the ordinal α, as follows:

- $R(p, \varphi, \psi, 2) \geq 0$ for any satisfiable p;
- for λ a limit ordinal, $R(p, \varphi, \psi, 2) \geq \lambda$ if $R(p, \varphi, \psi, 2) \geq \alpha$ for all $\alpha < \lambda$;
- $R(p, \varphi, \psi, 2) \geq \alpha + 1$ if there are satisfiable extensions p_1, p_2 of p and $b \in \mathcal{U}$ such that $\varphi(x, b) = 0$ is in p_1, $\psi(x, b) = 0$ is in p_2, $R(p_1, \varphi, \psi, 2) \geq \alpha$, and $R(p_2, \varphi, \psi, 2) \geq \alpha$.

We write $R(p, \varphi, \psi, 2) = \infty$ if $R(p, \varphi, \psi, 2) \geq \alpha$ for all ordinals α. Otherwise there is an ordinal γ such that $R(p, \varphi, \psi, 2) \geq \alpha$ holds iff $\alpha \leq \gamma$; in that case we set $R(p, \varphi, \psi, 2) = \gamma$. By compactness, if $R(p, \varphi, \psi, 2) \geq \omega$, then $R(p, \varphi, \psi, 2) = \infty$, so the values of R lie in $\mathbb{N} \cup \{\infty\}$.

14.3 Proposition. *Let $\mathcal{U}$ be a κ-universal domain for T. Then T is stable with respect to the discrete metric if and only if for all pairs of conditions $\varphi(x, y) = 0$, $\psi(x, y) = 0$ that are contradictory in $\mathcal{U}$, we have $R(\{d(x, x) = 0\}, \varphi, \psi, 2) < \omega$.*

Proof The argument is similar to the proof of the corresponding result in classical first-order model theory. See Proposition 2.2 in [3] for this proof in the cat framework. □

Next we introduce a notion of stability based on measuring the size of $S_1(T_A)$ by its density character with respect to the d-metric.

14.4 Definition. We say that T is λ-*stable* if for any $\mathcal{M} \models T$ and $A \subseteq M$ of cardinality $\leq \lambda$, there is a subset of $S_1(T_A)$ of cardinality $\leq \lambda$ that is dense in $S_1(T_A)$ with respect to the d-metric. We say that T is *stable* if T is λ-stable for some infinite λ.

14.5 Remark. There are two reasons why we have chosen to associate *stability* most closely with topological properties of the type spaces $S_1(T_A)$ expressed in terms of the d-metric. First, as will be seen in later sections, many theories of interest turn out to be ω-stable in this sense (but not necessarily ω-stable with respect to other natural topologies on $S_1(T_A)$ including the discrete topology). The second (and main) reason for this choice is that theories in continuous first-order logic that are λ-stable in this sense (*i.e.*, with respect to the d-metric) have properties analogous to those of λ-stable theories in classical first-order logic. (See

[3], [6].) For example ω-stable theories have prime models and ω-stable theories are λ-stable for all infinite λ. (See Remark 14.8 below.)

If we drop the quantitative aspect (*i.e.*, we drop λ) then the distinction between these two notions of stability disappears:

14.6 Theorem. *A theory T is stable if and only if T is stable with respect to the discrete metric.*

Proof Clearly if T is λ-stable with respect to the discrete metric then T is λ-stable.

Assume now that T is not stable with respect to the discrete metric. Let $\mathcal{U}$ be a κ-universal domain for T. By Proposition 14.3 there are formulas $\varphi_0(x,y)$, $\varphi_1(x,y)$ and $\epsilon > 0$ such that

$$\{\varphi_0(x,y) = 0, \varphi_1(x',y) = 0, d(x,x') \leq \epsilon\}$$

is not satisfiable in $\mathcal{U}$, and

$$R(\{d(x,x) = 0\}, \varphi_0(x,y), \varphi_1(x,y), 2) = \infty.$$

Given an infinite cardinal λ, let μ be a cardinal such that $2^{<\mu} \leq \lambda < 2^\mu$. We may assume that $\kappa > \lambda$. By the definition of the rank and the saturation of $\mathcal{U}$ we can find a sequence $\{b_\sigma \mid \sigma \in 2^{<\mu}\}$ of elements in $\mathcal{U}$ such that $\{\varphi_{\sigma(\alpha)}(x, b_{\sigma\restriction\alpha}) = 0 \mid \alpha < \mu\}$ is realized by some c_σ in $\mathcal{U}$, for every $\sigma \in 2^\mu$. Let $B = \{b_\sigma \mid \sigma \in 2^{<\mu}\}$. Then $\mathrm{card}(B)$ is $\leq \lambda$, yet $d(\mathrm{tp}(c_\sigma/B), \mathrm{tp}(c_\tau/B)) \geq \epsilon$ for all distinct $\sigma, \tau \in 2^\mu$. Therefore T is not λ-stable. $\square$

For many theories there are several natural topologies on type spaces that can be used as the basis of alternative notions of "λ-stability". This approach was considered by Iovino in [28, 29] for metric structures based on Banach spaces; he proved a generalization of Theorem 14.6 for such notions, in the setting of positive bounded formulas.

We continue the discussion of these type-counting notions of stability by quoting the Stability Spectrum Theorem. The proof is very much like the one in classical first-order logic. See Theorem 4.12 in [4] for this proof in the cat framework.

14.7 Theorem. *Let T be a stable theory and let $\mu(T)$ be the first cardinal in which T is stable. Then there exists a cardinal $\kappa = \kappa(T)$ such that T is λ-stable if and only if $\lambda = \mu(T) + \lambda^{<\kappa(T)}$.*

14.8 Remark. The previous result yields that any ω-stable theory is λ-stable for all infinite λ, because $\kappa(T)$ must be $\aleph_0$ for such theories. Most of the examples we treat in later sections are ω-stable.

14.9 Definition. A theory T is *superstable* if $\kappa(T) = \aleph_0$.

By the remark above, any ω-stable theory is superstable, but the converse implication is not true. For example, the theory of infinite dimensional Hilbert spaces with a generic automorphism is superstable but not ω-stable. (See [7].) There is a characterization of superstability in terms of the stability spectrum:

14.10 Proposition. *T is superstable if and only if it is λ-stable for all $\lambda \geq 2^{|T|}$.*

Proof See [6]. □

Now we begin discussing our second approach to stability, which is based on non-dividing:

14.11 Definition. Let $\mathcal{U}$ be a κ-universal domain for T and let B, C be small subsets of $\mathcal{U}$. Let $p(X, B)$ be a partial type over B in a possibly infinite tuple of variables X (so $p(X, Y)$ is a partial type without parameters). We say that $p(X, B)$ *divides* over C if there exists a C-indiscernible sequence $(B_i \mid i < \omega)$ in $\mathrm{tp}(B/C)$ such that $\bigcup_{i<\omega} p(X, B_i)$ is inconsistent with T.
Furthermore, if A, B, C are small sets in $\mathcal{U}$ such that $\mathrm{tp}(A/BC)$ does not divide over C, then we say that A *is independent from* B *over* C and we write $A \mathop{\smile\!\!\!\!\!\!\mid}_C B$.

This notion of independence has good properties in every stable theory, just as it does in classical first-order model theory:

14.12 Theorem. *Let $\mathcal{U}$ be a κ-universal domain for T. If T is stable, then the independence relation $\mathop{\smile\!\!\!\!\!\!\mid}$ defined using non-dividing satisfies the following properties (here A, B, etc., are any small subsets of $\mathcal{U}$ and M is a small elementary submodel of $\mathcal{U}$):*

1. *Invariance under automorphisms of $\mathcal{U}$.*
2. *Symmetry: $A \mathop{\smile\!\!\!\!\!\!\mid}_C B \Longleftrightarrow B \mathop{\smile\!\!\!\!\!\!\mid}_C A$.*
3. *Transitivity: $A \mathop{\smile\!\!\!\!\!\!\mid}_C BD$ if and only if $A \mathop{\smile\!\!\!\!\!\!\mid}_C B$ and $A \mathop{\smile\!\!\!\!\!\!\mid}_{B\cup C} D$.*

(4) *Finite Character:* $A \mathop{\smile\!\!\!\!\!\!|}_C B$ *if and only if* $a \mathop{\smile\!\!\!\!\!\!|}_C B$ *for all finite tuples* a *from* A.
(5) *Extension: For all* A, B, C *there exists* A' *such that* $A' \mathop{\smile\!\!\!\!\!\!|}_C B$ *and* $\mathrm{tp}(A/C) = \mathrm{tp}(A'/C)$.
(6) *Local Character: If* a *is any finite tuple, then there is* $B_0 \subseteq B$ *of cardinality* $\leq |T|$ *such that* $a \mathop{\smile\!\!\!\!\!\!|}_{B_0} B$.
(7) *Stationarity of types: If* $\mathrm{tp}(A/M) = \mathrm{tp}(A'/M)$, $A \mathop{\smile\!\!\!\!\!\!|}_M B$, *and* $A' \mathop{\smile\!\!\!\!\!\!|}_M B$, *then* $\mathrm{tp}(A/B \cup M) = \mathrm{tp}(A'/B \cup M)$.

Proof The proof follows the ideas of the corresponding result in classical first-order model theory. See, for example, [39].

A similar result has been proved in the more general framework of cats. Most of the proof can be found in [3, Theorems 1.51,2.8]), except that the extension property may fail in the setting considered there. In [4, Theorem 1.15] it is shown that in a cat that is *thick* and stable (more generally, if it is thick and simple), every type has a non-dividing extension. This applies to the setting considered in this paper, since any complete theory of metric structures gives rise to a Hausdorff cat, and every Hausdorff cat is thick. □

Stability can be characterized by the existence of an independence relation with suitable properties, just as in the classical first-order setting.

14.13 Definition. Let $\mathcal{U}$ be a κ-universal domain for T. A relation satisfying properties (1)–(7) in Theorem 14.12 is called a *stable independence relation* on $\mathcal{U}$.

14.14 Theorem. *Let* $\mathcal{U}$ *be a* κ*-universal domain for* T. *If* T *is stable, there is precisely one stable independence relation on* $\mathcal{U}$. *Moreover, if there exists a stable independence relation* $A \mathop{\smile\!\!\!\!\!\!|}^*_C B$ *on triples of small subsets of* $\mathcal{U}$, *then* T *is stable.*

Proof If T is a stable theory, the existence of a stable independence relation on $\mathcal{U}$ is given by Theorem 14.12. The rest of this Theorem is proved as in the classical first-order case. See, for example, [39]. □

Theorem 14.14 will play an important role in our treatment of application areas in sections 15–18. As is often true, the theories treated in those sections are already equipped with natural independence relations; for example, one has *orthogonality* in Hilbert spaces and *probabilistic independence* in probability spaces. (See sections 15 and 16 below.) Showing that such a relation is a stable independence relation is often not hard. Once this has been done, Theorem 14.14 implies that the theory in question is stable and that the natural candidate is indeed *the* relation of model-theoretic independence on $\mathcal{U}$ (yielding a complete understanding of non-dividing).

A third (equivalent) approach to stability is given by the notion of definability of types, which plays a central role in stability theory:

14.15 Definition. We say that a complete type $p \in S_n(B)$ is *definable* over a set A if for each formula $\varphi(x, y)$ with $x = x_1, \ldots, x_n$, there exists a predicate $\Psi(y)$ definable over A such that for all suitable tuples b in B we have that $\Psi(b)$ is the unique $r \in [0, 1]$ such that the condition $\varphi(x, b) = r$ is in p.

The following theorem characterizes stability in terms of definability of types:

14.16 Theorem. *The theory T is stable if and only if every type over a model $\mathcal{M}$ of T is definable over M.*

Proof Similar to the classical proof for first-order theories, but requiring a slightly more delicate analysis. See Theorem 8.5 in [6]. □

We conclude this section with the following result, which gives an alternative proof for Theorem 14.6 and which provides additional information about the stability spectrum.

14.17 Corollary. *Suppose T is stable. Then T is λ-stable with respect to the discrete metric for all λ that satisfy $\lambda = \lambda^{|T|}$.*

Proof Assume T is stable and that $\lambda = \lambda^{|T|}$. Let $\mathcal{M}$ be a model of T of density character λ; then M has cardinality at most $\lambda^{\aleph_0}$, which equals λ because of our special assumptions. Note that the number of definitions of types over M is at most $\lambda^{|T|} = \lambda$. Therefore, Theorem 14.16 yields that $S_1(T_M)$ has cardinality at most λ, as desired. □

The study of stability theory and its applications in analysis started with the work of Krivine and Maurey [32] around quantifier free formulas in Banach spaces. A study of stability and ω-stability for metric structures based on normed spaces is carried out in [28, 29, 30]. Theorem 14.6 was first proved (in a more general form) in Corollary 7.2 and

Corollary 7.3 in [28, page 88]. The proof provided here comes from [4, Remark 4.11]. A more general approach to stability in the setting of cats can be found in [3].

Stability of a general function (in particular, of a formula in continuous logic, *i.e.*, *local* stability) was introduced in [34] in the setting of functional analysis. Local stability of a formula in continuous logic is developed in [6].

There are several known examples of stable theories in the setting of metric structures, such as the theories of Hilbert spaces, probability spaces and L^p spaces. (See the next three sections.) Expansions of Hilbert spaces and probability spaces by generic automorphisms also turn out to be stable. (See [7, 8, 10] and the last section of this article.)

15 Hilbert spaces

A pre-Hilbert space H over $\mathbb{R}$ is a vector space over $\mathbb{R}$ with an inner product $\langle\ \rangle$ that satisfies the following properties:

(1) $\langle rx + sy, z\rangle = r\langle x, z\rangle + s\langle y, z\rangle$ for all $x, y, z \in H$ and all $r, s \in \mathbb{R}$;
(2) $\langle x, y\rangle = \langle y, x\rangle$ for all $x, y \in H$;
(3) $\langle x, x\rangle \in (0, \infty)$ for all nonzero $x \in H$; $\langle 0, 0\rangle = 0$.

On each pre-Hilbert space we define a norm by $\|x\| = \sqrt{\langle x, x\rangle}$. A pre-Hilbert space that is complete with respect to this norm is called a *Hilbert space*.

Let H be a pre-Hilbert space over $\mathbb{R}$. We treat H in continuous logic by identifying it with the many-sorted metric prestructure

$$\mathcal{M}(H) = \big((B_n(H) \mid n \geq 1), 0, \{I_{mn}\}_{m<n}, \{\lambda_r\}_{r\in\mathbb{R}}, +, -, \langle\ \rangle\big)$$

where $B_n(H) = \{x \in H \mid \|x\| \leq n\}$ for $n \geq 1$; 0 is the zero vector in $B_1(H)$; for $m < n$, $I_{mn} \colon B_m \to B_n$ is the inclusion map; for $r \in \mathbb{R}$ and $n \geq 1$, $\lambda_r \colon B_n(H) \to B_{nk}(H)$ is scalar multiplication by r, with k the unique integer satisfying $k \geq 1$ and $k - 1 \leq |r| < k$; furthermore, $+, - \colon B_n(H) \times B_n(H) \to B_{2n}(H)$ are vector addition and subtraction and $\langle\ \rangle \colon B_n(H) \to [-n^2, n^2]$ is the inner product, for each $n \geq 1$. The metric on each sort is given by $d(x, y) = \|x - y\|$.

There is an obvious continuous signature L such that for each pre-Hilbert space H, the many-sorted prestructure described above is an L-prestructure; the necessary bounds and moduli of uniform convergence are easy to specify. We see easily that if H is a pre-Hilbert space, then the completion of the L-prestructure $\mathcal{M}(H)$ is equal to $\mathcal{M}(\overline{H})$, where $\overline{H}$ is the Hilbert space obtained by completing H.

It is not difficult to show that there is an L-theory, which we will denote by HS, such that $\mathcal{M} \models HS$ if and only if there is a Hilbert space H such that $\mathcal{M} \cong \mathcal{M}(H)$. (One way to verify this is using Proposition 5.14; it is well known and easy to check that the class of L-structures isomorphic to some $\mathcal{M}(H)$, with H a Hilbert space, is closed under ultraproducts and under elementary substructures. It is also not difficult to extract a set of axioms for this class by directly translating the requirements that $\mathcal{M}$ comes from a vector space over $\mathbb{R}$ and that the inner product satisfies (1),(2),(3) above.)

For any Hilbert space H, we see that H is infinite dimensional if and only if $\mathcal{M}(H)$ satisfies the conditions

$$\big(\inf_{x_1}\dots\inf_{x_n}\max_{1\leq i,j\leq n}(|\langle x_i,x_j\rangle-\delta_{ij}|)\big)=0$$

for $n \geq 1$; here $\delta_{ij} = 1$ if $i = j$ and $\delta_{ij} = 0$ if $i \neq j$, and the variables $x_1, \dots, x_n$ range over the sort $B_1(H)$. Let IHS be the L-theory obtained by adding this infinite set of conditions to HS. An L-structure $\mathcal{M}$ is a model of IHS if and only if $\mathcal{M}$ is isomorphic to $\mathcal{M}(H)$ for some infinite dimensional Hilbert space H. In what follows we will identify H with $\mathcal{M}(H)$ when applying model theoretic techniques and concepts, such as types.

Note that IHS is κ-categorical for every infinite cardinal κ. Therefore IHS is a complete theory. In the rest of this section we will show that IHS admits quantifier elimination and is ω-stable.

Let $x, y \in H$ and let $A \subset H$. By $\overline{A}$ we mean the norm closure of the linear span of A. Let $P_{\bar{A}}(x)$ be the projection of x on the subspace $\overline{A}$. We denote by $A^\perp$ the set $\{z \in H : \langle a, z\rangle = 0 \text{ for all } a \in A\}$; it is a closed subspace of H known as the *orthogonal complement* of $\overline{A}$, since H is the Hilbert space direct sum of $\overline{A}$ and $A^\perp$.

15.1 Lemma. *Let H be an infinite dimensional Hilbert space, with $c_1, \dots, c_n, d_1, \dots, d_n \in H$. Then $(c_1, \dots, c_n)$ and $(d_1, \dots, d_n)$ realize the same type over $A \subset H$ if and only if $P_{\bar{A}}(c_i) = P_{\bar{A}}(d_i)$ and $\langle c_i, c_j\rangle = \langle d_i, d_j\rangle$ for all $1 \leq i, j \leq n$.*

Proof If $\mathrm{tp}(c_1, \dots, c_n/A) = \mathrm{tp}(d_1, \dots, d_n/A)$ then $\langle c_i, c_j\rangle = \langle d_i, d_j\rangle$ for $i, j \leq n$ and for every $a, b \in A$, $\langle c_i - b, a\rangle = \langle d_i - b, a\rangle$; thus $P_{\bar{A}}(c_i) = P_{\bar{A}}(d_i)$.

Conversely, assume that $P_{\bar{A}}(c_i) = P_{\bar{A}}(d_i)$ and $\langle c_i, c_j\rangle = \langle d_i, d_j\rangle$ for $i, j \le n$. Then $c_i - P_{\bar{A}}(c_i), d_i - P_{\bar{A}}(d_i) \in A^\perp$ and

$$\langle c_i - P_{\bar{A}}(c_i), c_j - P_{\bar{A}}(c_j)\rangle = \langle d_i - P_{\bar{A}}(d_i), d_j - P_{\bar{A}}(d_j)\rangle.$$

Using the Gram-Schmidt process we can build an automorphism of H that fixes $\bar{A}$ pointwise and takes $c_i - P_{\bar{A}}(c_i)$ to $d_i - P_{\bar{A}}(d_i)$ for all $1 \le i \le n$. □

15.2 Corollary. *The theory IHS admits quantifier elimination.*

Proof We apply Proposition 13.6. Suppose $\mathcal{M}_1, \mathcal{M}_2 \models IHS$; let $\mathcal{N}$ be a substructure of M_1 and $f : \mathcal{N} \to \mathcal{M}_2$ an embedding. Let $\mathcal{M}_2' \succeq \mathcal{M}_2$ be such that the orthogonal complement of M_2 in M_2' has dimension $\ge \dim(M_1)$. We can extend f so it maps an orthonormal basis of $M_1 \cap N^\perp$ into an orthonormal subset of $M_2' \cap M_2^\perp$ and then extend f linearly to all of M_1. By Lemma 15.1, such a map is an embedding. □

15.3 Lemma. *Let H be an infinite dimensional Hilbert space and let $A \subset H$. Then the definable closure of A equals $\bar{A}$.*

Proof By passing to an elementary extension of H, which does not change $\mathrm{dcl}(A)$, we may assume that $\bar{A}$ is a proper subspace of H.

We first show that if $c \in \bar{A}$, then $c \in \mathrm{dcl}(A)$. Given $c \in \bar{A}$, there is a Cauchy sequence $\{c_n : n \ge 1\}$ of elements in the space spanned by A such that $\lim_{n\to\infty} c_n = c$. We may assume that $\|c_n - c\| \le 1/(2n)$ for $n \ge 1$. Let $\varphi_n(x) = \|x - c_n\| \dot- 1/(2n)$. Then the family of formulas $\{\varphi_n(x) \mid n \ge 1\}$ and numbers $\{\delta_n = 1/(2n) \mid n \ge 1\}$ shows that $\{c\}$ is A-definable.

Assume now that $c \notin \bar{A}$, so $c - P_{\bar{A}}(c) \neq 0$. Take any $y \in A^\perp$ such that $\|y\| = \|c - P_{\bar{A}}(c)\|$. Then $\mathrm{tp}(c/A) = \mathrm{tp}(P_{\bar{A}}(c) + y/A)$. Since $A^\perp$ is not the 0 subspace, this shows that $\mathrm{tp}(c/A)$ has realizations in H that are different from c, and thus $c \notin \mathrm{dcl}(A)$. □

15.4 Proposition. *Let H be an infinite dimensional Hilbert space. For each $x, y \in H$ and $A \subset H$ we have*

$$d(\mathrm{tp}(x/A), \mathrm{tp}(y/A))^2 = \|P_{\bar{A}}(x) - P_{\bar{A}}(y)\|^2 + \big|\|x - P_{\bar{A}}(x)\| - \|y - P_{\bar{A}}(y)\|\big|^2$$

Proof Let $x, y \in H$ and let $A \subset H$. If $\mathrm{tp}(x'/A) = \mathrm{tp}(x/A)$ and

$\mathrm{tp}(y'/A) = \mathrm{tp}(y/A)$, then

$$\begin{aligned}\|x' - y'\|^2 &= \|P_{\bar{A}}(x') - P_{\bar{A}}(y')\|^2 + \|(x' - P_{\bar{A}}(x')) - (y' - P_{\bar{A}}(y'))\|^2 \\ &\geq \|P_{\bar{A}}(x) - P_{\bar{A}}(y)\|^2 + \big|\|x - P_{\bar{A}}(x)\| - \|y - P_{\bar{A}}(y)\|\big|^2.\end{aligned}$$

For the reverse inequality, let $x_\perp = x - P_{\bar{A}}(x)$ and $y_\perp = y - P_{\bar{A}}(y)$. If $x_\perp = 0$ the result is clear, so we may assume that $x_\perp \neq 0$. Let $\alpha = \|y_\perp\|/\|x_\perp\|$ and let $z = \alpha x_\perp$. Then we have $\mathrm{tp}(y/A) = \mathrm{tp}(P_{\bar{A}}(y) + z/A)$ by Lemma 15.1 and

$$\begin{aligned}\|x - (P_A(y) + z)\|^2 &= \|P_{\bar{A}}(x) - P_{\bar{A}}(y)\|^2 + \big|\|x_\perp - \alpha x_\perp\|\big|^2 \\ &= \|P_{\bar{A}}(x) - P_{\bar{A}}(y)\|^2 + \big|\|x_\perp\| - \|y_\perp\|\big|^2\end{aligned}$$

by the Pythagorean theorem. □

15.5 Proposition. *The theory IHS is ω-stable.*

Proof Let H be an infinite dimensional Hilbert space and let $A \subset H$ be countable. Let $T_A = \mathrm{Th}(\mathcal{M}(H), a)_{a \in A}$. We may assume that the dimension of $A^\perp$ (in H) is ω, so H is separable. Using Lemma 15.1 it is easy to show that every 1-type over A is realized in H. Therefore $S_1(A)$ is separable with respect to the d-metric. □

We close this section by giving a concrete description in terms of familiar Hilbert space concepts of the independence relation that is associated to the theory IHS. As usual, this gives an alternate proof that IHS is stable, although it does not identify the values of λ for which IHS is λ-stable.

In what follows, we fix a cardinal number $\kappa > 2^{\aleph_0}$ and we fix a κ-universal domain H for the theory IHS.

15.6 Definition. Whenever A, B, C are small subsets of H, we write $A \mathop{\smile\!\!\!\!\!\!\;\;|}^*_C B$ to mean that $P_{\bar{C}}(a) = P_{\overline{C \cup B}}(a)$ for all $a \in A$.

15.7 Lemma. *Let $A, B, C \subset H$ be small. Then $A \mathop{\smile\!\!\!\!\!\!\;\;|}^*_C B$ if and only if $a - P_{\bar{C}}(a) \perp b - P_{\bar{C}}(b)$ for all $a \in A$ and $b \in B$.*

Proof Assume first that $A \mathop{\smile\!\!\!\!\!\!\;\;|}^*_C B$, so for all $a \in A$, $P_{\bar{C}}(a) = P_{\overline{C \cup B}}(a)$. Then $a - P_{\bar{C}}(a) \in (C \cup B)^\perp$, so $a - P_{\bar{C}}(a) \perp b - P_{\bar{C}}(b)$ for any $b \in B$. Now assume that $a - P_{\bar{C}}(a) \perp b - P_{\bar{C}}(b)$ for all $a \in A$ and $b \in B$. Then $a - P_{\bar{C}}(a) \perp b$ for all $a \in A$ and $b \in B$ and thus $P_{\bar{C}}(a) = P_{\overline{C \cup B}}(a)$ for all $a \in A$. □

15.8 Theorem. *The relation $\mathop{\smile\!\!\!\!\!\!|}^*$ is a stable independence relation on H. Therefore, $\mathop{\smile\!\!\!\!\!\!|}^*$ is identical to the independence relation $\mathop{\smile\!\!\!\!\!\!|}$ based on non-dividing for the stable theory IHS.*

Proof We prove directly that $\mathop{\smile\!\!\!\!\!\!|}^*$ has all seven properties of a stable independence relation:

1. (1) Invariance: Let $f \in \mathrm{Aut}(H)$. Then for every $u, v \in H$, $\langle u|v\rangle = \langle f(u)|f(v)\rangle$. Therefore, if E is any subspace of H, it is easy to see that f carries P_E to $P_{f(E)}$ and carries $E^\perp$ to $f(E)^\perp$. This makes it clear (from the definition) that $f(A) \mathop{\smile\!\!\!\!\!\!|}^*_{f(C)} f(B)$ is equivalent to $A \mathop{\smile\!\!\!\!\!\!|}^*_C B$ for any small subsets A, B, C of H.
2. (2) Symmetry: This follows from Lemma 15.7.
3. (3) Transitivity: This follows from the definition.
4. (4) Finite character: This is immediate from the definition.
5. (5) Extension: By finite character and compactness, it suffices to prove the property for finite tuples. Let $a_1, \dots, a_n \in H$ and let $B, C \subseteq H$ be small. Since $B \cup C$ is small, $(B \cup C)^\perp$ is an infinite dimensional subspace of H. Hence there are $c_1, \dots, c_n \in (C \cup B)^\perp$ such that $\mathrm{tp}(c_1, \dots, c_n) = \mathrm{tp}(a_1 - P_{\bar{C}}(a_1), \dots, a_n - P_{\bar{C}}(a_n))$. Let $a_i' = P_{\bar{C}}(a_i) + c_i$, for $i = 1, \dots, n$. Then
$$\mathrm{tp}(a_1', \dots, a_n'/C) = \mathrm{tp}(a_1, \dots, a_n/C)$$
and $\{a_1', \dots, a_n'\} \mathop{\smile\!\!\!\!\!\!|}^*_C B$.
6. (6) Local Character: Given $a_1, \dots, a_n$ and B there exists a countable subset B_0 of B such that each of $P_{\bar{B}}(a_i)$ is an element of $\overline{B_0}$. Then we have $a_1, \dots, a_n \mathop{\smile\!\!\!\!\!\!|}^*_{B_0} B$.
7. (7) Stationarity of types: We will show that the property holds for general sets, that is, we do not need to assume that the underlying set C is an elementary substructure of H. By finite character, it suffices to prove the property when A is a finite tuple. So let $a = (a_1, \dots, a_n), a' = (a_1', \dots, a_n') \in H^n$ and let $C, B \subseteq H$ be small. Assume that $\mathrm{tp}(a/C) = \mathrm{tp}(a'/C)$ and that $a \mathop{\smile\!\!\!\!\!\!|}^*_C B$, $a' \mathop{\smile\!\!\!\!\!\!|}^*_C B$. Then for every $i = 1, \dots, n$,
$$P_{\overline{B \cup C}}(a_i) = P_{\bar{C}}(a_i) = P_{\bar{C}}(a_i') = P_{\overline{B \cup C}}(a_i').$$
Thus $\mathrm{tp}(a/B \cup C) = \mathrm{tp}(a'/B \cup C)$.

The second statement follows from the first by Theorem 14.14. □

Note that the proof of Theorem 15.8 shows that all types of tuples in IHS are stationary.

Model theoretic studies of infinite dimensional Hilbert spaces are carried out in [9] and in [1]. A direct proof to characterize non-dividing in terms of orthogonality can be found in Corollary 2 and Lemma 8 in [9]. Proposition 15.5 appears in [28].

16 Probability spaces

In this section we give an introduction to the model theory of probability spaces using their measure algebras as metric structures.

A *probability space* is a triple $(X, \mathcal{B}, \mu)$, where X is a set, $\mathcal{B}$ is a σ-algebra of subsets of X and μ is a σ-additive measure on $\mathcal{B}$ such that $\mu(X) = 1$.

We say that $B \in \mathcal{B}$ is an *atom* if $\mu(B) > 0$ and there does not exist any $B' \in \mathcal{B}$ satisfying $B' \subseteq B$ and $0 < \mu(B') < \mu(B)$. Further, $B \in \mathcal{B}$ is *atomless* if there is no atom that is a subset of B. The probability space $(X, \mathcal{B}, \mu)$ is *atomless* if X itself is atomless. It is well known that if B is an atomless element of $\mathcal{B}$ and $0 \leq r \leq 1$ then there exists $B' \in \mathcal{B}$ satisfying $B' \subseteq B$ and $\mu(B') = r \cdot \mu(B)$. (See [20, Section 41] for a discussion.) This uniformity in the property of being atomless plays a role in axiomatizing the property in continuous logic. (See below.)

We write $A_1 \sim_\mu A_2$, and say that $A_1, A_2 \in \mathcal{B}$ determine the same *event*, if the symmetric difference of the sets, denoted by $A_1 \triangle A_2$, has measure zero. Clearly $\sim_\mu$ is an equivalence relation. We denote the class of $A \in \mathcal{B}$ under the equivalence relation $\sim_\mu$ by $[A]_\mu$. The collection of equivalence classes of $\mathcal{B}$ modulo $\sim_\mu$ is called the set of *events* and it is denoted by $\widehat{\mathcal{B}}$. The operations of complement, union and intersection are well defined for events and make $\widehat{\mathcal{B}}$ a σ-algebra; in addition, μ induces a well-defined, strictly positive, countably additive probability measure on $\widehat{\mathcal{B}}$. We refer to $\widehat{\mathcal{B}}$ as the *measure algebra* and to $(\widehat{\mathcal{B}}, \mu)$ as the *measured algebra* associated to $(X, \mathcal{B}, \mu)$.

Given $(X, \mathcal{B}, \mu)$, we build a 1-sorted metric structure (called a *probability structure*)

$$\mathcal{M} = (\widehat{\mathcal{B}}, 0, 1, \cdot^c, \cap, \cup, \mu)$$

whose metric is given by $d([A]_\mu, [B]_\mu) = \mu(A \triangle B)$. Here 0 is the event of measure zero, 1 the event of measure one, and $\cdot^c, \cap, \cup$ are the Boolean operations induced on $\widehat{\mathcal{B}}$. The modulus of uniform continuity for $\cdot^c$ is the identity $\Delta(\epsilon) = \epsilon$ and the moduli of uniform continuity for $\cup$ and $\cap$ are given by $\Delta'(\epsilon) = \epsilon/2$. We sometimes write a^{-1} for a^c and a^{+1} for a, when a is an element of $\widehat{\mathcal{B}}$.

Let L be the signature associated to these probability structures. The following L-conditions are true in all probability structures. Indeed, Theorem 16.1 shows they axiomatize that class of structures.

(1) Boolean algebra axioms:
Each of the usual axioms for Boolean algebras is the closure under universal quantifiers of an equation between terms (see [27, page 38]) and thus it can be expressed in continuous logic as a condition. For example, the axiom $\forall x \forall y (x \cup y = y \cup x)$ is equivalent to $\sup_x \sup_y \big(d(x \cup y, y \cup x)\big) = 0$.

(2) Measure axioms:
$\mu(0) = 0$ and $\mu(1) = 1$;
$\sup_x \sup_y \big(\mu(x \cap y) \dot- \mu(x)\big) = 0$;
$\sup_x \sup_y \big(\mu(x) \dot- \mu(x \cup y)\big) = 0$;
$\sup_x \sup_y |(\mu(x) \dot- \mu(x \cap y)) - (\mu(x \cup y) \dot- \mu(y))| = 0$.
The last three axioms express that $\mu(x \cup y) + \mu(x \cap y) = \mu(y) + \mu(x)$ for all x, y.

(3) Connection between d and μ:
$\sup_x \sup_y |d(x, y) - \mu(x \Delta y)| = 0$ where $x \Delta y$ denotes the Boolean term giving the symmetric difference: $(x \cap y^c) \cup (x^c \cap y)$.

We denote the set of L-conditions above by PrA.

16.1 Theorem. *Let $\mathcal{M}$ be an L-structure with underlying metric space (M, d). Then $\mathcal{M}$ is a model of PrA if and only if $\mathcal{M}$ is the probability structure associated to a probability space $(X, \mathcal{B}, \mu)$ as above.*

Proof It is clear that probability structures satisfy the conditions in PrA. A proof that such structures are metrically complete is given in [38, Chapter 7]. For the converse, let $\mathcal{M} \models PrA$; recall this implies that underlying metric space of $\mathcal{M}$ is complete. In [38, Chapter 7] it is discussed how to realize $\mathcal{M}$ as a probability structure. $\square$

To say that a model is *atomless* we need the following axiom; it states that every set of positive measure can be cut nearly "in half" measurably. As noted above, this is well known to hold in a probability space if and only if the space is atomless.

(4) $\sup_x \inf_y |\mu(x \cap y) - \mu(x \cap y^c)| = 0$.

We denote by APA the set of axioms PrA together with (4). Its models are exactly the probability structures obtained from atomless probability spaces.

16.2 Proposition. *The theory APA is separably categorical (and therefore APA is complete).*

Proof Let $\mathcal{M} \models APA$ be separable. As is shown in [38, Chapter 7], $\mathcal{M}$ is the probability structure associated to a countably generated probability space that is necessarily atomless. A familiar back and forth argument shows that any two such probability spaces are isomorphic in a measure preserving manner. □

Next we characterize the d-metric on spaces of types for APA. To do this, we need the following special case of the Radon-Nikodym theorem:

16.3 Theorem. *(Radon-Nikodym; see [16, Theorem 3.8]) Let $(X, \mathcal{B}, \mu)$ be a probability space, let $\mathcal{C} \subseteq \mathcal{B}$ be a σ-subalgebra and let $A \in \mathcal{B}$. Let a be the event corresponding to A. Then there is a unique $g_a \in L^1(X, \mathcal{C}, \mu)$ such that for any $B \in \mathcal{C}$, $\int_B g_a d\mu = \int_B \chi_A d\mu$. Such an element g_a is called the* conditional probability of a with respect to $\mathcal{C}$ *and it is denoted by $\mathbb{P}(a|\mathcal{C})$.*

The next lemma provides an explicit form for the d-metric on probability structures (this formula was known to analysts [37, Lemma 6.3] in the case $C = \emptyset$).

16.4 Lemma. *Let $\mathcal{M} \models APA$ be a κ-universal domain, with $\kappa \geq \omega_1$. Assume that $\mathcal{M}$ is the probability structure associated to the probability space $(Y, \mathcal{D}, m)$. Let $a = (a_1, \dots, a_n) \in M^n$, $b = (b_1, \dots, b_n) \in M^n$ be partitions of the probability structure. Let $C \subseteq M$ be small, and let $\mathcal{C}$ be the σ-subalgebra of $\mathcal{D}$ generated by the measurable sets A such that the event of A is in C. Then*

$$d(\mathrm{tp}(a/C), \mathrm{tp}(b/C)) = \max_{1 \leq i \leq n} \|\mathbb{P}(a_i|\mathcal{C}) - \mathbb{P}(b_i|\mathcal{C})\|_1$$

where $\| \ \|_1$ is the L_1-norm.

Proof This is Lemma 3.14 in [10]. □

16.5 Corollary. *(Ben Yaacov [2]) Let $\mathcal{M} \models APA$ be a κ-universal domain, with $\kappa \geq \omega_1$, and let $C \subseteq M$ be small. Let $\mathcal{C}$ be obtained from C as in the previous result. Let $a = (a_1, \dots, a_n) \in M^n$, $b = (b_1, \dots, b_n) \in M^n$ be arbitrary. Then $\mathrm{tp}(a/C) = \mathrm{tp}(b/C)$ iff*

$$\mathbb{P}(a_1^{i_1} \cap \dots \cap a_n^{i_n}|\mathcal{C}) = \mathbb{P}(b_1^{i_1} \cap \dots \cap b_n^{i_n}|\mathcal{C})$$

for all n-tuples $(i_1, \ldots, i_n)$ *from* $\{+1, -1\}$.

Proof Apply the previous result to the tuples of atoms in the finite subalgebras generated by $a_1, \ldots, a_n$ and $b_1, \ldots, b_n$ respectively. □

16.6 Proposition. *The theory APA admits quantifier elimination.*

Proof If $\mathcal{M} \models APA$ and $a_1, \ldots, a_n \in M$, then the previous result shows that $\text{tp}(a_1, \ldots, a_n)$ is determined by the measures of the atoms of the finite Boolean subalgebra generated by $a_1, \ldots, a_n$. This together with Proposition 13.2 implies that APA admits quantifier elimination. □

16.7 Proposition. *Let* $\mathcal{M}$ *be a model of APA and* $C \subseteq M$. *The* definable closure of C *is the smallest* σ*-algebra of events containing* C.

Proof This follows from quantifier elimination plus an analysis of restricted quantifier free formulas. □

Non-dividing in probability structures has a natural characterization:

16.8 Proposition. *(Ben Yaacov [2, Theorem 2.10]) Let* $\mathcal{M} \models APA$ *be a* κ*-universal domain, with* $\kappa \geq \omega_1$*, and let* $C \subseteq M$ *be small. Assume that* $\mathcal{M}$ *is the probability structure associated to the probability space* $(X, \mathcal{B}, \mu)$. *Let* $a_1, \ldots, a_n; b_1, \ldots, b_m \in M$. *Let* $\mathcal{C}$ *be the* σ*-subalgebra of* $\mathcal{B}$ *generated by a collection of measurable sets whose events make up* C *and let* $\mathcal{C}_b$ *be the* σ*-subalgebra of* $\mathcal{B}$ *generated by a larger collection of measurable sets whose events make up* $C \cup \{b_1, \ldots, b_m\}$. *Then* $\text{tp}(a_1, \ldots, a_n/C \cup \{b_1, \ldots, b_m\})$ *does not divide over* C *if and only if*

$$\mathbb{P}(a_1^{i_1} \wedge \cdots \wedge a_n^{i_n} | \mathcal{C}_b) = \mathbb{P}(a_1^{i_1} \wedge \cdots \wedge a_n^{i_n} | \mathcal{C})$$

for all n-tuples $(i_1, \ldots, i_n)$ *from* $\{-1, 1\}$.

16.9 Proposition. *(1) The theory APA is* ω*-stable.*
(2) Let $\mathcal{N} \models APA$ *and consider* $a \in N^n$ *and* $C \subseteq N$. *Then* $\text{tp}(a/C)$ *is stationary.*

Proof Let $\mathcal{M} \models APA$ be a κ-universal domain, with $\kappa \geq \omega_1$. Let $(X, \mathcal{B}, \mu)$ be a probability space whose events correspond to the elements of M. For part (1), let $C \subseteq M$ be countable. We may assume that C is closed under finite intersections, unions and complements. Let $\mathcal{C}$ be a set of measurable sets whose set of events is C. We may choose $\mathcal{C}$ so that it is a countable Boolean algebra. Let $\langle\mathcal{C}\rangle$ be the σ-algebra generated by $\mathcal{C}$, let Step(C) be the set of step functions in $L^1(X, \mathcal{C}, m)$ with coefficients in $\mathbb{Q}$ and let $\mathcal{F} = \{\mathrm{tp}(a/C) \mid \mathbb{P}(a|\langle\mathcal{C}\rangle) \in \mathrm{Step}(C)\}$. Then $\mathcal{F}$ is a countable set of types. It follows from Lemma 16.4 that $\mathcal{F}$ is a dense subset of the space of 1-types over C with respect to the d-metric. Therefore APA is ω-stable. The proof of (2) follows from 16.8. □

A model theoretic study of (atomless) probability spaces is carried out in [2]. The author shows that they give rise to a compact abstract theory (see [1] for the definition) and that, in this cat, the notion of independence obtained from non-dividing (Proposition 16.8) agrees with probabilistic independence. A characterization of the d-metric (Lemma 16.4) is derived in [10], which also contains the proof of ω-stability given above. In general, the presentation of the material in this section follows [10] closely.

17 L^p Banach lattices

Let X be a set, U a σ-algebra on X and μ a σ-additive measure on U, and let $p \in [1, \infty)$. We denote by $L^p(X, U, \mu)$ the space of (equivalence classes of) U-measurable functions $f \colon X \to \mathbb{R}$ such that $\|f\| = (\int |f|^p d\mu)^{1/p} < \infty$. We consider this space as a Banach lattice (complete normed vector lattice) over $\mathbb{R}$ in the usual way; in particular, the lattice operations $\wedge, \vee$ are given by pointwise maximum and minimum.

We will work on models of the form

$$((B_n \mid n \geq 1), 0, \{I_{mn}\}_{m<n}, \{\lambda_r\}_{r\in\mathbb{R}}, +, -, \wedge, \vee, \| \ \|)$$

where $B_n = B_n(L^p(X, U, \mu)) = \{f \in L^p(X, U, \mu) \mid \|f\| \leq n\}$ and $I_{mn} \colon B_m \to B_n$ is the inclusion map for $m < n$. The metric on each B_n is given by $d(f, g) = \|f - g\|$. The diameter of B_n is $2n$ and the values of the predicate $\| \ \|$ on B_n are in $[0, n]$. The operations $+, -, \wedge, \vee$ map B_n into B_{2n}. For $r \in \mathbb{R}$ with $k - 1 < |r| \leq k$, where $k \geq 1$ is an integer, the operation λ_r (of scalar multiplication by r) maps B_n into B_{kn}.

The moduli of uniform continuity for the norm and for the inclusion maps I_{mn} are all given by $\Delta(\epsilon) = \epsilon$. The moduli of uniform continuity for $+, -, \wedge, \vee$ are all given by $\Delta'(\epsilon) = \epsilon/2$. For $r \in \mathbb{R}$ with $k-1 < |r| \leq k$, where k is an integer ≥ 1, the modulus of uniform continuity of λ_r is given by $\Delta_{\lambda_r}(\epsilon) = \epsilon/k$.

Let L denote the signature just described.

Basic analysis and probability

We start with a review of some results from analysis that we use to approach L^p spaces model theoretically.

A measure space (X, U, μ) is called *decomposable* (also called *strictly localizable*) if there exists a partition $\{X_i \mid i \in I\} \subseteq U$ of X into measurable sets such that $\mu(X_i) < \infty$ for all $i \in I$ and such that for any subset A of X, $A \in U$ iff $A \cap X_i \in U$ for all $i \in I$ and, in that case, $\mu(A) = \sum_{i \in I} \mu(A \cap X_i)$.

17.1 Convention. Throughout this section we require that all measure spaces are decomposable.

There is no loss of generality in adopting this convention: the representation theorem for abstract L^p spaces shows that every such space (and, in particular therefore, every concrete $L^p(X, U, \mu)$ space) can be represented in this way with (X, U, μ) being decomposable. (See the proof of Theorem 3 in [11], for example.)

Let E be any Banach lattice and $f \in E$. The *positive part* of f is $f \vee 0$, and it is denoted f^+. The *negative part* of f is $f^- = (-f)^+$, and one has $f = f^+ - f^-$ and the *absolute value* of f is given by $|f| = f^+ + f^-$. Further, f is *positive* if $f = f^+$ and f is *negative* if $-f$ is positive. For $f, g \in E$, one has $f \geq g$ iff $f - g$ is positive.

Let (X, U, μ) be a measure space. A measurable set $S \in U$ is an *atom* if $\mu(S) > 0$ and there does not exist any $S' \in U$ satisfying $S' \subseteq S$ and $0 < \mu(S') < \mu(S)$. One calls (X, U, μ) *atomless* if it has no atoms.

If E is a Banach lattice and $x \in E$, a *component* of x is $y \in E$ such that $|y| \wedge |x - y| = 0$. If (X, U, μ) is a measure space and $E = L^p(X, U, \mu)$, then the components of $x \in E$ are the results of restricting x to some measurable subset of the support of x.

17.2 Notation. Let (X, U, μ) and (Y, V, μ) be a measure spaces. We write $(Y, V, \mu) \subseteq (X, U, \mu)$ to mean that $Y \in U$ and $V \subseteq U$.

Model theory of L^p Banach lattices

In this section, unless stated otherwise, we work on the unit ball. So all elements under consideration and all quantifiers range over B_1.

It is routine to write down L-conditions expressing the following axioms, which are true in $L^p(X, U, \mu)$, where (X, U, μ) is a measure space.

(1) The Banach lattice axioms, described in [35, pages 47–49, 81].

(2) Axioms for abstract L^p spaces, which state that

$$\|x \wedge y\|^p \leq \|x\|^p + \|y\|^p \leq \|x + y\|^p$$

whenever x and y are positive.

We write LpL for the theory axiomatized above.

17.3 Proposition. *(Axiomatizability) Let $\mathcal{M}$ be an L-structure with underlying metric space (M, d). Then $\mathcal{M}$ is a model of LpL if and only if there is a measure space (X, U, μ) such that $\mathcal{M}$ is isomorphic to $L^p(X, U, \mu)$.*

Proof See the proof of Theorem 3 in [11], for example. □

To ensure that the measure space representing the structure under consideration is atomless we need an additional axiom:

(3) $\sup_x \inf_y \big(\max(\big|\|y\| - \|x^+ - y\|\big|, \|y \wedge (x^+ - y)\|\big) = 0$.

It is obvious that this condition is satisfied in any $L^p(X, U, \mu)$ for which (X, U, μ) is atomless. For the converse, note that (3) states that for every positive function $u = x^+$ and every $\epsilon > 0$ there is some y such that $\|y\|$ and $\|u - y\|$ differ by at most ϵ and $\|y \wedge (u - y)\| \leq \epsilon$. Assume $u \neq 0$. By taking ϵ small enough and subtracting $y \wedge (u - y)$ from y, we get a nontrivial component of u. Hence, in models of LpL that also satisfy condition (3), there are no atoms.

We denote by $ALpL$ the set of conditions $LpL + (3)$. We just proved:

17.4 Proposition. *If $\mathcal{M}$ is an L-structure, then $\mathcal{M}$ is a model of the theory $ALpL$ if and only if there is an atomless measure space (X, U, μ) such that $\mathcal{M}$ is isomorphic to $L^p(X, U, \mu)$.*

17.5 Fact. (Quantifier elimination) Let $\mathcal{M}$ be the L^p space of an atomless measure space and let $a, b \in M^n$. If a and b have the same quantifier free type in $\mathcal{M}$, then a and b have the same type in $\mathcal{M}$. That is, if $\|t(a)\| = \|t(b)\|$ in $\mathcal{M}$ for every term $t(x_1, \ldots, x_n)$, then a and b have the same type in $\mathcal{M}$. From this observation and Proposition 13.6 it follows that $ALpL$ admits quantifier elimination.

For a proof of Fact 17.5, see Example 13.18 in [24].

Note that Fact 17.5 is not true without the assumption that $\mathcal{U}$ is atomless; atoms and non-atoms can have the same quantifier free type but never have the same type.

17.6 Fact. (Separable categoricity) If $\mathcal{M} \models ALpL$ be separable, then $\mathcal{M}$ is isomorphic to $L^p([0,1], \mathcal{B}, m)$, where $\mathcal{B}$ is the σ-algebra of Lebesgue measurable sets and m is Lebesgue measure.

Note that ω-categoricity need not be preserved when we add constants to the language:

17.7 Example. Let f and g be any two norm 1, positive elements of $L^p([0,1], \mathcal{B}, m)$. By Fact 17.5, we get

$$(L^p([0,1], \mathcal{B}, m), f) \equiv (L^p([0,1], \mathcal{B}, m), g).$$

However, there are two possible isomorphism types of such structures, depending on whether or not the support of the adjoined function has measure 1 or not.

17.8 Fact. Let $\mathcal{M}$ be a model of $ALpL$ and $A \subseteq M$. The definable closure of A in $\mathcal{M}$ is the closed linear sublattice of $\mathcal{M}$ generated by A.

Proof This follows from quantifier elimination plus an analysis of restricted quantifier free formulas. □

17.9 Remark. Let $\mathcal{M}$ be the L^p space of a measure space and let C be a closed linear sublattice of $\mathcal{M}$. One can use the representation theorem for abstract L^p spaces (see [33, pages 15–16], part(2) of the axiomatization) to show that there exist measure spaces (X, U, μ) and (Y, V, μ) satisfying $(Y, V, \mu) \subseteq (X, U, \mu)$, as well as an isomorphism Φ from $L^p(X, U, \mu)$ onto $\mathcal{M}$ such that Φ maps $L^p(Y, V, \mu)$ exactly onto C.

The previous remark has interesting consequences. Let $\mathcal{M}$ and C be as in Remark 17.9 and let $f \in M^n$. In [5] it is proved that the type over C realized by f in $\mathcal{M}$ is characterized by the joint conditional distribution of f over the σ-algebra associated to C (V in the notation above). The proof of this result is beyond the scope of this paper. However the ideas behind it are illustrated by the special case of a single characteristic function:

17.10 Proposition. *Let $\mathcal{M}$ and C be as in Remark 17.9. Suppose $(Y, V, \mu) \subset (X, U, \mu)$ are measure spaces such that $C = L^p(Y, V, \mu)$ and $M = L^p(X, U, \mu)$. Let $A, B \subset Y$ be such that $\chi_A, \chi_B \in M$. Then $\mathbb{P}(A|V) = \mathbb{P}(B|V)$ implies $\mathrm{tp}(\chi_A/C) = \mathrm{tp}(\chi_B/C)$.*

Proof Assume that $\mathbb{P}(A|V) = \mathbb{P}(B|V)$. By quantifier elimination, to show that $\mathrm{tp}(\chi_A/C) = \mathrm{tp}(\chi_B/C)$ it suffices to prove that for any $g \in C^l$ and any lattice term $t(x, y)$, we have $\|t(\chi_A, g)\|^p = \|t(\chi_B, g)\|^p$.

Let ν be the measure on Borel subsets D of $\mathbb{R}^{1+l}$ defined by $\nu(D) = \mu\{x \in X : (\chi_A, g)(x) \in D\}$. Since $\mathbb{P}(A|V) = \mathbb{P}(B|V)$, we have that $\nu(D) = \mu\{x \in X : (\chi_B, g)(x) \in D\}$ for any Borel $D \subset \mathbb{R}^{1+l}$. Then, by the change of variable formula,

$$\begin{aligned} \int_X |t(\chi_A(x), g(x))|^p d\mu(x) &= \int_{\mathbb{R}^{1+l}} |t(r, s)|^p d\nu(r, s) \\ &= \int_X |t(\chi_B(x), g(x))|^p d\mu(x) \end{aligned}$$

and hence $\|t(\chi_A, g)\|^p = \|t(\chi_B, g)\|^p$ as desired. □

Finally we show stability.

17.11 Theorem. *(Henson [22]) The theory ALpL is ω-stable.*

Proof Let $\mathcal{U}$ be a κ-universal domain for $ALpL$, with $\kappa \geq \omega_1$, and let $A \subseteq \mathcal{U}$ be countably infinite. Then $\mathrm{dcl}(A)$ is a closed linear sublattice of $\mathcal{U}$ and thus we can find measure spaces $(Y, V, \mu) \subseteq (X, U, \mu)$ such that $\mathrm{dcl}(A) = L^p(Y, V, \mu)$ and $\mathcal{U} = L^p(X, U, \mu)$. Let $T_A = \mathrm{Th}((\mathcal{U}, a)_{a\in A})$.

Any function $h \in \mathcal{U}$ can be written as $f + g$, where the support of f is contained in Y and the support of g is disjoint from Y. Moreover, $\mathrm{tp}((f + g)/A)$ is determined by $\mathrm{tp}(f/A)$ and $\mathrm{tp}(g/A)$.

Therefore, to find the density character of $S_1(T_A)$ it suffices to consider the following two cases:

Let $f \in \mathcal{U}$ be an element supported on Y. It suffices to consider the case where f is a simple function, since every function is a limit of simple functions. We identify $L^p(Y, V, \mu)$ with its canonical image in the space $L^p((Y, V, \mu) \otimes ([0, 1], \mathcal{B}, m))$, where $([0, 1], \mathcal{B}, m)$ is the standard Lebesgue space. Since $\mathcal{U}$ is sufficiently saturated, we may assume that $L^p((Y, V, \mu) \otimes ([0, 1], \mathcal{B}, m))$ is a closed linear sublattice of $\mathcal{U}$. Using quantifier elimination (see Fact 17.5), the fact that f is a simple function, and Proposition 17.10, we can find $f' \in L^p((Y, V, \mu) \otimes ([0, 1], \mathcal{B}, m))$ such that $\mathrm{tp}(f/A) = \mathrm{tp}(f'/A)$. Since $L^p((Y, V, \mu) \otimes ([0, 1], \mathcal{B}, m))$ is separable, the density character of the space of types of functions supported on Y is ω.

Let $g \in \mathcal{U}$ be an element whose support is disjoint from Y. The type $\mathrm{tp}(g/A)$ is determined by $\|g^+\|$ and $\|g^-\|$. Let $B, C \in U$ be disjoint from Y, each of measure one. The set $\{\mathrm{tp}(c_1\chi_B - c_2\chi_C) \mid c_1, c_2 \in \mathbb{Q}^+\}$

is a countable dense subset of the space of types of functions disjoint to $\mathrm{dcl}(A)$. □

The basic model-theoretic properties of L^p spaces, such as axiomatizability and quantifier elimination (see Fact 17.5) are proved in [24]. Further results about the model theory of these structures are given in [5], including a characterization of the d-metric and an analysis of non-dividing in terms of conditional expectation, as well as stability theoretic properties such as stationarity of types over sets.

18 Probability spaces with generic automorphism

In this section we will study the existentially closed structures of the form $(\mathcal{M}, \tau)$, where $\mathcal{M}$ is a probability structure and τ is an automorphism of $\mathcal{M}$. We show this class is axiomatizable and its theory is stable. We also discuss the model-theoretic meaning of some results in ergodic theory.

Lebesgue spaces and their automorphisms

There are two approaches to isomorphisms on probability spaces. On the one hand, we have measure preserving point maps between the spaces; on the other, we have measure preserving maps between measured algebras.

18.1 Definition. Let $(X_1, \mathcal{B}_1, \mu_1)$, $(X_2, \mathcal{B}_2, \mu_2)$ be probability spaces and let $\widehat{\mathcal{B}}_1$, $\widehat{\mathcal{B}}_2$ be their measure algebras. By an *isomorphism* of the measured algebras we mean a bijection $\Phi\colon \widehat{\mathcal{B}}_1 \to \widehat{\mathcal{B}}_2$ that preserves complements, countable unions and intersections and satisfies $\mu_2(\Phi(b)) = \mu_1(b)$ for all $b \in \widehat{\mathcal{B}}_1$. The probability spaces are said to be *conjugate* if their measured algebras are isomorphic.

18.2 Definition. Let $(X_1, \mathcal{B}_1, \mu_1)$, $(X_2, \mathcal{B}_2, \mu_2)$ be probability spaces and let $\widehat{\mathcal{B}}_1$, $\widehat{\mathcal{B}}_2$ be their measure algebras. Let $C_1 \in \mathcal{B}_1$, $C_2 \in \mathcal{B}_2$ with $\mu_1(C_1) = 1 = \mu_2(C_2)$. An invertible measure preserving transformation $\Phi\colon C_1 \to C_2$ is called an *isomorphism* between $(X_1, \mathcal{B}_1, \mu_1)$ and $(X_2, \mathcal{B}_2, \mu_2)$. If $(X_1, \mathcal{B}_1, \mu_1) = (X_2, \mathcal{B}_2, \mu_2)$, we call Φ an *automorphism.* For $b \in \widehat{\mathcal{B}}_1$, let $B \in \mathcal{B}_1$ be such that $[B]_{\mu_1} = b$. Let $\widehat{\Phi}(b) = [\Phi(B \cap C_1)]_{\mu_2}$. The induced map $\widehat{\Phi}\colon \widehat{B}_1 \to \widehat{B}_2$ is an isomorphism and it is called an *induced isomorphism* of the measured algebras.

Clearly any two isomorphic probability spaces are conjugate; however, the converse does not hold in general. The next definition concerns a well-known special class of probability spaces where the converse does hold.

18.3 Definition. A probability space $(I, \mathcal{L}, m)$ is a *Lebesgue space* if it is isomorphic to a probability space that is the disjoint union of two spaces:

(1) One that is a countable (or finite) set of points $\{y_1, y_2, \dots\}$, each of positive measure.
(2) The space $([0, s], \mathcal{L}([0, s]), l)$, where $\mathcal{L}([0, s])$ is the Lebesgue σ-algebra on $[0, s]$ and l is Lebesgue measure.

Here $s = 1 - \sum_{i=1}^{\infty} p_i$, where $p_i > 0$ is the measure of $\{y_i\}$.

On Lebesgue spaces the notion of isomorphism and conjugacy coincide. (See Theorem 2.2 in [40].)

18.4 Definition. Let $(Y, \mathcal{C}, \mu)$ be an atomless probability space and let τ_Y be an automorphism of $(Y, \mathcal{C}, \mu)$. We say that τ_Y is *aperiodic* if for every $n \in \mathbb{N}^+$, the set $\{y \in Y \mid \tau_Y^n(y) = y\}$ has measure zero.

For the rest of this section we will study aperiodic maps and their properties. A good source for this material is the book of Halmos [21, pages 69–76] on ergodic theory. One of the key tools for studying aperiodic automorphisms is:

18.5 Theorem. *(Rokhlin's Lemma [21, page 71]) Let $(Y, \mathcal{C}, \mu)$ be an atomless probability space and τ_Y an aperiodic automorphism of this space. Then for every positive integer n and $\epsilon > 0$, there exists a measurable set $E \in \mathcal{C}$ such that the sets $E, \tau_Y(E), \dots, \tau_Y^{n-1}(E)$ are disjoint and $\mu(\cup_{i<n} \tau_Y^i(E)) > 1 - \epsilon$.*

18.6 Remark. Let $(Y, \mathcal{C}, \mu)$ be an atomless probability space and let τ_Y be an automorphism of this space. Let $\mathcal{N}$ be the probability structure induced by $(Y, \mathcal{C}, \mu)$ and let τ be the automorphism of $\mathcal{N}$ induced by τ_Y. Then τ_Y is aperiodic iff

$$\inf_e \max\big(|1/n - \mu(e)|, \mu(e \cap \tau(e)), \mu(e \cap \tau^2(e)), \dots, \mu(e \cap \tau^{n-1}(e))\big) = 0$$

for all $n \geq 1$.

For the rest of this section we fix an atomless Lebesgue space $(I, \mathcal{L}, m)$.

18.7 Fact. [21, page 74] Let $A, B \in \mathcal{L}$ be such that $m(A) = m(B)$. Then there is an automorphism η of $(I, \mathcal{L}, m)$ such that $m(\eta(A) \triangle B) = 0$.

From now on, G denotes the group of measure preserving automorphisms on $(I, \mathcal{L}, m)$, where we identify two maps if they agree on a set of measure one. There is a natural representation of G in $\mathbb{B}(L^2(I, \mathcal{L}, m))$ (the space of bounded linear operators on $L^2(I, \mathcal{L}, m)$); it sends $\tau \in G$ to the unitary operator U_τ defined for all $f \in L^2(I, \mathcal{L}, m)$ by $U_\tau(f) = f \circ \tau$. The norm topology on $\mathbb{B}(L^2(I, \mathcal{L}, m))$ pulls back to a group topology on G, which is called the *uniform topology* on G in [21, page 69]. For $\tau, \eta \in G$, let $\rho(\tau, \eta) = m(\{x \in X \mid \tau(x) \neq \eta(x)\})$. It is shown in [21, pages 72–73] that ρ is a metric for the uniform topology.

18.8 Definition. We call a map $\eta \in G$ a *cycle* of period n if there is a set $E \in \mathcal{L}$ of measure $1/n$ such that $E, \eta(E), \dots, \eta^{n-1}(E)$ are pairwise disjoint, and $\eta^n = id \restriction_X$ (up to measure zero).

18.9 Remark. (1) Let $\tau \in G$ be aperiodic. For every $n > 0$ there is a cycle $\eta \in G$ of period n such that $\rho(\tau, \eta) \leq 2/n$. (By Rokhlin's Lemma.)
(2) Given any two cycles $\eta_1, \eta_2 \in G$ of period n, there is $\gamma \in G$ such that $\gamma^{-1}\eta_1\gamma = \eta_2$. (This follows from Fact 18.7.)
(3) Let $\tau_1, \tau_2 \in G$ be aperiodic. Then for every $\epsilon > 0$, there is $\gamma \in G$ such that $\rho(\tau_1, \gamma^{-1}\tau_2\gamma) \leq \epsilon$. (This follows from (1) and (2).)

Existentially closed structures

We denote by L be the language of probability structures and by APA the theory of atomless probability structures. Write L_τ for the language L expanded by a unary function with symbol τ and let APA_τ be the theory $APA \cup$ "τ is an automorphism". We can axiomatize APA_τ by adding to APA the following conditions:

(1) $\sup_x \left|\mu(x) - \mu(\tau(x))\right| = 0$
(2) $\sup_x \inf_y \left|\mu(x \triangle \tau(y))\right| = 0$
(3) $\sup_x \sup_y \left|\mu(\tau(x \cup y) \triangle (\tau(x) \cup \tau(y)))\right| = 0$
(4) $\sup_x \sup_y \left|\mu(\tau(x \cap y) \triangle (\tau(x) \cap \tau(y)))\right| = 0$

(Note that (1) expresses that τ is measure preserving, (2) that τ is surjective, and (3,4) that τ is a Boolean homomorphism.)

We write tp for types in the language L and tp_τ for types in the language L_τ. Let $\mathcal{M}$ be the probability structure associated to $(I, \mathcal{L}, m)$.

Recall that G is the group of automorphisms of $(I, \mathcal{L}, m)$, where we identify two maps if they agree on a set of measure one.

18.10 Proposition. *Let τ_I, τ_I' be aperiodic automorphisms of $(\mathcal{I}, L, m)$ and let τ, τ' be the corresponding induced automorphisms of $\mathcal{M}$. Then $(\mathcal{M}, \tau) \equiv (\mathcal{M}, \tau')$.*

Proof An application of Remark 18.9(3) with the values for ϵ ranging over the sequence $\{1/n \mid n \in \mathbb{N}^+\}$ together with the uniform continuity of formulas shows that $(\mathcal{M}, \tau) \equiv (\mathcal{M}, \tau')$. □

Thus any two models induced by aperiodic transformations on an atomless *Lebesgue space* have the same elementary theory. The aim of this section is to study this theory.

18.11 Definition. Let $(X, \mathcal{B}, \mu)$ be a probability space, let $\widehat{\mathcal{B}}$ be the corresponding measure algebra of events and let τ be an automorphism of the measure algebra $\widehat{\mathcal{B}}$. The map τ is called *aperiodic* if it satisfies the following condition corresponding to the conclusion of Rokhlin's lemma:

$$\inf_e \max\big(|1/n - \mu(e)|, \mu(e \cap \tau(e)), \mu(e \cap \tau^2(e)), \ldots, \mu(e \cap \tau^{n-1}(e))\big) = 0$$

for all $n \geq 1$.

Let $APAA$ be the theory APA_τ together with the conditions in L_τ describing that τ is an aperiodic automorphism.

18.12 Lemma. *The theory $APAA$ is complete.*

Proof Let (M_1, τ_1) and (M_2, τ_2) be two models of $APAA$. Then there are separable models $(M_i', \tau_i') \models APAA$ that are elementarily equivalent to (M_i, τ_i) for $i = 1, 2$ respectively. By separable categoricity of APA, for $i = 1, 2$ we may assume that M_i' is the probability structure associated to a Lebesgue space and that τ_i' is induced by an automorphism η_i of the corresponding Lebesgue space. Then η_1 and η_2 are aperiodic automorphisms of the Lebesgue spaces and thus by Proposition 18.10 we have $(M_1', \tau_1') \equiv (M_2', \tau_2')$. □

18.13 Remark. The theory of aperiodic automorphisms on probability structures is the limit, as n goes to infinity, of the theory of a probability structure formed by n atoms $\{a_1, \ldots, a_n\}$ each of measure $1/n$, equipped with a cycle τ of period n; that is, τ is a permutation of $\{a_1, \ldots, a_n\}$ such that

$$\{a_1, \tau(a_1), \ldots, \tau^{n-1}(a_1)\} = \{a_1, \ldots, a_n\}.$$

Indeed, for each sentence φ in the language of $APAA$, let $v(n, \varphi)$ denote the value φ gets in the n-point probability space with a cycle of period n. On the other hand, let $v(\varphi)$ be the value φ gets in any model of $APAA$. (This value is well defined since $APAA$ is complete.) Then, $\lim_{n\to\infty} v(n, \varphi) = v(\varphi)$. Otherwise there would be a non-principal ultrafilter $\mathcal{U}$ on the positive integers such that if $\mathcal{M}$ is the $\mathcal{U}$-ultraproduct of the family consisting of the n-point probability structures with cycles of period n, then $\varphi^{\mathcal{M}} \neq v(\varphi)$. But this is a contradiction, because any such $\mathcal{M}$ is a model of $APAA$.

Our next aim is to show that $APAA$ is model complete. The techniques used for the proof are similar to the ones used in proving the completeness of the theory $APAA$, but now we need to include parameters.

18.14 Proposition. *The theory APAA is model complete. That is, if* $(N_0, \tau_0) \subseteq (N_1, \tau_1)$ *are models of APAA, then* $(N_0, \tau_0) \preceq (N_1, \tau_1)$.

Proof Let $(\mathcal{M}, \tau) \subseteq (\mathcal{N}, \tau)$ be separable models of $APAA$. We may assume that there is an atomless Lebesgue space $(X, \mathcal{B}, m)$ such that $\mathcal{M}$ is the model induced by $(X, \mathcal{B}, m)$ and that there is an atomless Lebesgue space $(Y, \mathcal{C}, m)$ such that $\mathcal{N}$ is the structure induced by $(Y, \mathcal{C}, m)$. Furthermore, we may assume that there is an aperiodic automorphism τ_X on $(X, \mathcal{B}, m)$ that induces the action of τ on $\mathcal{M}$ and that there is an aperiodic automorphism τ_Y on $(Y, \mathcal{C}, m)$ that induces the action of τ on $\mathcal{N}$. Note that both maps τ_X and τ_Y are aperiodic.

Let $a_1, \ldots, a_p \in M$. Take $A_1^X, \ldots, A_p^X \in \mathcal{B}$ such that $[A_j^X]_m = a_j$ for $j = 1, \ldots, p$, and $A_1^Y, \ldots, A_p^Y \in \mathcal{C}$ such that $[A_j^Y]_m = a_j$ for $j = 1, \ldots, p$.

Let $\epsilon > 0$ and let n be a positive integer such that $1/n < \epsilon$. By Rokhlin's Lemma, there is $B \in \mathcal{B}$ such that $B, \tau_X(B), \ldots, \tau_X^{n-1}(B)$ are disjoint and $m(\cup_{i<n} \tau_X^i(B)) \geq 1 - \epsilon$. Let P^X be the partition of B generated by $\tau_X^{-i}(\tau_X^i(B) \cap A_j^X)$ for $1 \leq j \leq p$ and $0 \leq i < n$. Since $(\mathcal{M}, \tau) \subseteq (\mathcal{N}, \tau)$, there is $C \in \mathcal{C}$ such that $[C]_m = [B]_m$. Let P^Y be the partition of C generated by $\tau_Y^{-i}(\tau_Y^i(C) \cap A_j^Y)$ for $1 \leq j \leq p$ and $0 \leq i < n$.

Since $(X, \mathcal{B}, m)$ and $(Y, \mathcal{C}, m)$ are Lebesgue spaces, by Fact 18.7 there is $h_0\colon B \to C$ a measure preserving bijection such that

$$h_0(\tau_X^{-i}(\tau_X^i(B) \cap A_j^X)) = \tau_Y^{-i}(\tau_Y^i(C) \cap A_j^Y)$$

(up to measure zero) for $1 \le j \le p$ and $0 \le i < n$. Extend h_0 to $\cup_{1<i<n}\tau_X^i(B)$ by setting $h_0(\tau_X^i(x)) = \tau_Y^i(h_0(x))$ for $x \in B$, $0 < i < n$.

Let $Z^X = X \setminus \cup_{1<i<n}\tau_X^i(B)$ and let $Z^Y = Y \setminus \cup_{1<i<n}\tau_Y^i(C)$. Extend h_0 to a measure preserving bijection h from $(X, \mathcal{B}, m)$ to $(Y, \mathcal{C}, m)$ by defining $h\colon Z^X \to Z^Y$ so that $h(Z^X \cap A_j^X) = Z^Y \cap A_j^Y$ for all $1 \le j \le p$.

Note that h induces a map $\widehat{h}$ from $\mathcal{M}$ to $\mathcal{N}$ satisfying $\widehat{h}(a_j) = a_j$ for all j and that $\rho(\tau^X, h^{-1}\tau^Y h) \le 2\epsilon$. The proposition follows from the uniform continuity of formulas. □

18.15 Definition. We say that $(\mathcal{M}, \tau) \models APA_\tau$ is *existentially closed* if whenever $(\mathcal{N}, \tau) \models APA_\tau$, $(\mathcal{N}, \tau) \supseteq (\mathcal{M}, \tau)$, $a \in M^n$ and $\varphi(x, y)$ is a quantifier free formula such that $(\mathcal{N}, \tau) \models \inf_x \varphi(x, a) = 0$, then $(\mathcal{M}, \tau) \models \inf_x \varphi(x, a) = 0$. (Here x, y are disjoint finite sequences of distinct variables, with y of length n.)

18.16 Lemma. *Let $\mathcal{M}$ be an L_τ-structure. Then $\mathcal{M} \models APAA$ if and only if $\mathcal{M}$ is an existentially closed model of APA_τ.*

Proof Any model of APA_τ can be embedded in a model of $APAA$. Indeed, if $\mathcal{M} \models APA_\tau$, then $\mathcal{M}$ can be embedded in the product of $\mathcal{M}$ and the L_τ-structure that is based on the unit circle with normalized Lebesgue measure plus the rotation through an irrational multiple of π. It is easy to see that this product is a model of $APAA$. Since $APAA$ is axiomatized by adding a set of inf-conditions to APA_τ, this yields that any existentially closed model of APA_τ is a model of $APAA$. The other direction follows from the previous proposition. □

In ergodic theory, *joinings* (see [18, page 125]) give different ways of amalgamating two probability structures with automorphisms into a common extension. In particular, the *relative independent joining over a common factor* (described in [18, page 127]) corresponds to a free amalgamation of two models $\mathcal{M}_1$, $\mathcal{M}_2$ of APA_τ over a common substructure $\mathcal{N}$. Call this new structure $\mathcal{M}_1 \oplus_{\mathcal{N}} \mathcal{M}_2$.

18.17 Theorem. *The theory APAA has elimination of quantifiers.*

Proof We apply Proposition 13.6. Let $\mathcal{M}_1$, $\mathcal{M}_2 \models APAA$, let $\mathcal{N}$ be a substructure of $\mathcal{M}_1$ and let $f : \mathcal{N} \to \mathcal{M}_2$ be an embedding. We may assume that f is the identity. Since $APAA$ is model complete, $\mathcal{M}_1 \oplus_{\mathcal{N}} \mathcal{M}_2$ is an elementary extension of $\mathcal{M}_2$ and the canonical embedding of $\mathcal{M}_1$ into this space gives the desired extension of f. □

18.18 Proposition. *Let (M, τ) be a κ-universal domain of APAA, with $\kappa \geq \omega_1$, and let $A \subseteq M$. Then the definable closure of A in L_τ is the smallest σ-algebra of events containing $\tau^i(A)$ for all $i \in \mathbb{Z}$.*

Proof This follows from quantifier elimination plus an analysis of restricted quantifier free formulas. □

Independence and stability

In this section we introduce an abstract notion of independence, which we call τ-independence, and show that it agrees with non-dividing. This idea follows the approach used in [14, Section 3] to characterize non-dividing inside a first-order stable structure expanded by a generic automorphism. We reserve the word "independence" here for independence of events in the sense of probability structures. Fix a κ-universal domain $(\mathcal{U}, \tau) \models APAA$, where $\kappa \geq \omega_1$.

18.19 Definition. Let $A, B, C \subset \mathcal{U}$ be small. We write $A \mathop{\smile\!\!\!\!\!\!|}^{\tau}_{C} B$, and say that A *is τ-independent from B over C*, if $\mathrm{dcl}_\tau(A)$ is independent from $\mathrm{dcl}_\tau(B)$ over $\mathrm{dcl}_\tau(C)$.

We will show that $\mathop{\smile\!\!\!\!\!\!|}^{\tau}$ is a stable independence relation. We start by proving a strong form of stationarity:

18.20 Proposition. *Let $a, b \in \mathcal{U}^n$ and let $C \subseteq D \subseteq \mathcal{U}$. Suppose that $\mathrm{tp}_\tau(a/C) = \mathrm{tp}_\tau(b/C)$ and that $a \mathop{\smile\!\!\!\!\!\!|}^{\tau}_{C} D$ and $b \mathop{\smile\!\!\!\!\!\!|}^{\tau}_{C} D$. Then $\mathrm{tp}_\tau(a/D) = \mathrm{tp}_\tau(b/D)$.*

Proof Let a, b, C, D be as above. Then for every $k < \omega$,

$$\mathrm{tp}(\tau^{-k}(a), \dots, \tau^k(a))/\mathrm{dcl}_\tau(C)) = \mathrm{tp}(\tau^{-k}(b), \dots, \tau^k(b))/\mathrm{dcl}_\tau(C)).$$

By stationarity of types in models of APA, we get

$$\mathrm{tp}(\tau^{-k}(a), \dots, \tau^k(a))/\mathrm{dcl}_\tau(D)) = \mathrm{tp}(\tau^{-k}(b), \dots, \tau^k(b))/\mathrm{dcl}_\tau(D)).$$

Since this equality holds for all $k < \omega$, by quantifier elimination of $APAA$, $\mathrm{tp}_\tau(a/D) = \mathrm{tp}_\tau(b/D)$. □

18.21 Corollary. *The theory APAA is stable and the relation of τ-independence agrees with non-dividing.*

Proof We note that τ-independence is a stable independence relation on κ-universal domains for $APAA$, where $\kappa \geq \omega_1$. Using the properties of independence in a κ-universal domain $\mathcal{U}$ for APA, it is clear that τ-independence satisfies: invariance, symmetry, transitivity, extension, local character and finite character. By the previous proposition it also satisfies stationarity. Applying Theorem 14.14, this shows that τ-independence agrees with non-dividing and $APAA$ is stable. □

18.22 Remark. The theory $APAA$ is not ω-stable. For every irrational $\alpha \in [0, 1]$, consider the model of APAA that is based on the unit circle with normalized Lebesgue measure and is equipped with the aperiodic automorphism corresponding to rotation by the angle $2\pi\alpha$. Let p_α be the 1-type realized in this model by the event corresponding to the semicircle in the upper half plane.

Assume now that α, β are irrational and that $\alpha - \beta$ is also irrational. Let $a \models p_\alpha$ and let $b \models p_\beta$. For every $\epsilon > 0$ there is $n < \omega$ such that $d(a, \sigma^n(a)) = \mu(a \triangle \sigma^n(a)) < \epsilon$ while $d(b, \sigma^n(b)) > 1 - \epsilon$. It follows that $d(a, b) + d(\sigma^n(a), \sigma^n(b)) \geq 1 - 2\epsilon$. Since $d(a, b) = d(\sigma^n(a), \sigma^n(b))$, we conclude that $d(a, b) \geq \frac{1}{2} - \epsilon$. Therefore $d(p_\alpha, p_\beta) \geq \frac{1}{2}$.

Since we can choose $2^{\aleph_0}$ irrationals in $[0, 1]$, each two of which are linearly independent over $\mathbb{Q}$, this yields $2^{\aleph_0}$ types over $\emptyset$ each two of which have distance $\geq \frac{1}{2}$ from each other.

Further results about probability algebras with a generic automorphism can be found in [10]. In particular, it is shown there how *entropy* can be seen as a model-theoretic rank.

References

[1] Itaï Ben Yaacov, Positive model theory and compact abstract theories, *Journal of Mathematical Logic* **3**, 2003, 85–118.
[2] Itaï Ben Yaacov, Schroedinger's cat, *Israel Journal of Mathematics* **153**, 2006, 157–191.
[3] Itaï Ben Yaacov, Simplicity in compact abstract theories, *Journal of Mathematical Logic* **3**, 2003, 163–191.
[4] Itaï Ben Yaacov, Uncountable dense categoricity in cats, *The Journal of Symbolic Logic* **70**, 2005, 829-860.
[5] Itaï Ben Yaacov, Alexander Berenstein and C. Ward Henson, Model-theoretic independence in the Banach lattices $L^p(\mu)$, submitted.
[6] Itaï Ben Yaacov and Alex Usvyatsov, Continuous first order logic and local stability, submitted.
[7] Itaï Ben Yaacov, Alex Usvyatsov and Moshe Zadka, The theory of a Hilbert space with a generic automorphism, in preparation.
[8] Alexander Berenstein, Hilbert spaces with generic groups of automorphisms, *Archive for Mathematical Logic* **46**, 2007, 289–299.
[9] Alexander Berenstein and Steven Buechler, Homogeneous expansions of Hilbert spaces, *Annals of Pure and Applied Logic* **128**, 2004, 75–101.
[10] Alexander Berenstein and C. Ward Henson, Model theory of probability spaces with an automorphism, submitted.
[11] Jean Bretagnolle, Didier Dacunha-Castelle, Jean-Louis Krivine, Lois stables et espaces L^p, *Annals Inst. Henry Poincaré Sect. B (N.S.)* **2**, 1965/1966, 231–259.
[12] C. C. Chang and H. Jerome Keisler, *Continuous Model Theory*, Princeton University Press, 1966.
[13] C. C. Chang and H. Jerome Keisler. *Model theory*, volume 73 of Studies in Logic and the Foundations of Mathematics, North-Holland, third edition, 1990.
[14] Zoé Chatzidakis and Anand Pillay, Generic structures and simple theories, *Annals of Pure and Applied Logic* **95**, 1998, 71–92.
[15] Didier 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.
[16] Gerald B. Folland, *Real Analysis*, John Wiley and Sons, 1984.
[17] Leonard Gillman and Meyer Jerison, *Rings of Continuous Functions*, Van Nostrand, 1960.
[18] Eli Glasner, *Ergodic Theory via Joinings*, Mathematical Surveys and Monographs, 101, American Mathematical Society, 2003.
[19] Mikhail Gromov, *Metric Structures for Riemannian and Non-Riemannian Spaces*, Birkhäuser, 1999.
[20] Paul A. Halmos, *Measure Theory*, Van Nostrand, 1950.
[21] Paul A. Halmos, *Lectures on Ergodic Theory*, Chelsea, 1956.
[22] C. Ward Henson, Model theory of separable Banach spaces (abstract), *The Journal of Symbolic Logic* **52**, 1987, 1059–1060.
[23] C. Ward Henson, Nonstandard hulls of Banach spaces, *Israel Journal of Mathematics* **25**, 1976, 108–144.
[24] C. Ward Henson and José Iovino, Ultraproducts in analysis, in *Analysis and Logic*, London Mathematical Society Lecture Notes Series, vol. 262, 2002, 1–113.

[25] 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.
[26] C. Ward Henson and Hernando Tellez, Algebraic closure in continuous logic, *Revista Colombiana de Matemáticas*, to appear.
[27] Wilfrid Hodges, *Model Theory*, Cambridge University Press, 1993.
[28] José Iovino, Stable Banach space structures, I: Fundamentals, in *Models, algebras and proofs* (Bogotá, 1995), Lecture Notes in Pure and Applied Mathematics, 203, Marcel Dekker, 1999, 77–95.
[29] José Iovino, Stable Banach space structures, II: Forking and compact topologies, in *Models, algebras and proofs* (Bogotá, 1995), Lecture Notes in Pure and Applied Mathematics, 203, Marcel Dekker, 1999, 97–117.
[30] José Iovino, The Morley rank of a Banach space, *The Journal of Symbolic Logic* **61**, 1996, 85–118.
[31] H. Jerome Keisler, Ultraproducts and elementary classes, *Proceedings of the Koninklijke Nederlandse Akademie van Wetenschappen (Indagationes mathematicae)*, Series A, **23**, 1961, 477–495.
[32] Jean-Louis Krivine and Bernard Maurey, Espaces de Banach stables, *Israel Journal of Mathematics* **39**, 1981, 273–295.
[33] Joram Lindenstrauss and Lior Tzafriri, *Classical Banach Spaces, II; Function Spaces*, Springer Verlag, 1979.
[34] Yves Raynaud, Espaces de Banach superstables, distances stables et homéomorphismes uniformes, *Israel Journal of Mathematics* **44**, 1983, 33–52.
[35] Helmut H. Schaefer, *Banach Lattices and Positive Operators*, Springer Verlag, 1974.
[36] Saharon Shelah, Every two elementarily equivalent models have isomorphic ultrapowers, *Israel Journal of Mathematics* **10**, 1971, 224–233.
[37] Paul Shields, *The Theory of Bernoulli Shifts*, Chicago Lectures in Mathematics, The University of Chicago Press, 1973.
[38] D. A. Vladimirov, *Boolean Algebras in Analysis*, Kluwer Academic Publishers, 2002.
[39] Frank O. Wagner, *Simple Theories*, Kluwer Academic Publishers, 2000.
[40] Peter Walters, *An Introduction to Ergodic Theory*, Springer Verlag, 1982.