Metric spaces are universal for bi-interpretation with metric structures

James Hanson

September 8, 2019

Abstract

In the context of metric structures introduced by Ben Yaacov, Berenstein, Henson, and Usvyatsov[5], we exhibit an explicit encoding of metric structures in countable signatures as pure metric spaces in the empty signature, showing that such structures are universal for bi-interpretation among metric structures with positive diameter. This is analogous to the classical encoding of arbitrary discrete structures in finite signatures as graphs, but is stronger in certain ways and weaker in others. There are also certain fine grained topological concerns with no analog in the discrete setting.

1 Introduction

It is a well known fact[6] that any discrete structure with finite signature can be encoded as a graph in a particularly strong way:

Fact 1.1. *For any finite signature $\mathcal{L}$ there is a sentence χ in a language with a single binary predicate such that every model of χ is a graph and the class of models of χ is bi-interpretable with the class of $\mathcal{L}$-structures with more than one element. Furthermore this bi-interpretation preserves embeddings and is computable in the sense that presentations of models of χ are uniformly computable from presentations of the corresponding $\mathcal{L}$-structure and vice versa.*

This immediately implies that the set of tautologies involving a single binary predicate is undecidable, even though so called 'monadic first-order logic,' involving only unary predicates, is decidable. This is in contrast to the situation in continuous first-order logic, introduced in [5]. There is an easy encoding of a graph (V, E) as a metric space (V, d), where

$$d(x,y) = \begin{cases} 0 & x = y \\ \frac{1}{2} & xEy \\ 1 & \text{otherwise} \end{cases} .$$

So the set of continuous tautologies in the empty signature is undecidable for any reasonable notion of computable continuous formulas. Moreover, discrete structures can be encoded as metric spaces, in light of Fact 1.1.

The proof of Fact 1.1 uses a 'tag construction,' in which each tuple $x_0, x_1, ..., x_{k-1}$ related by some relation P is connected by a tag which is engineered to distinguish each x_i and to be distinguishable from tags corresponding to relations other than P. This particular construction does not generalize in any satisfactory way to metric structures, but the main result of this paper is a generalization of Fact 1.1:

Theorem 1.2. *For any countable metric signature $\mathcal{L}$ and $r > 0$, there is a theory T in the empty signature such that the class of models of T is bi-interpretable with the class of $\mathcal{L}$-structures with diameter $\geq r$. This bi-interpretation preserves embeddings and d-finiteness of types. If the original structure is not strongly infinite dimensional, then the interpreted structure will also not be strongly infinite dimensional. Furthermore, the bi-interpretation is computable in the sense that presentations of models of T are uniformly computable from presentations of the corresponding $\mathcal{L}$-structure and vice versa.*

There are some improvements in Theorem 1.2 over Fact 1.1, namely that the encoding works in the empty signature—which is largely cosmetic—and that we can encode countable signatures rather than just finite ones. 'd-finiteness of types' is a technical niceness condition introduced in [1] that will be discussed below. Strong infinite dimensionality is relevant from the point of view of computable structure theory, since the continuous degree of a point in a finite dimensional or weakly infinite dimensional metric space is always total[7]. These two concepts play no essential role in the construction.

The restriction that the metric structures have diameter uniformly bounded below is the necessary analog of the 'more than one element' restriction. A simple compactness argument shows that we could never have bi-interpretability between a single elementary class of metric spaces and the class of all $\mathcal{L}$-structures of positive diameter. In both the discrete case and the metric case we could avoid this non-uniformity by appending a new sort to every structure that always contains precisely two elements distance 1 apart. Also it should be noted that this is a non-issue from a computable structure theory point of view.

Finally there's the issue of finite axiomatizability, which the generalization loses, although, as will be discussed at the end of the paper in Section 6.1, there is no clear analog of finite axiomatizability in continuous logic.

2 Preliminaries

In the interest of notational brevity, we will describe one step of the bi-interpretation informally before defining the concept of a metric signature rigorously:

Fact 2.1. *Every many-sorted metric signature can be recast as a purely relational metric signature with $[0, 1]$-valued predicates and metrics.*

So from now on all predicate symbols will be $[0, 1]$-valued and in particular all sorts will have diameter ≤ 1.

There are some trivial subtleties if we allow ourselves predicates with zero-length ranges or other such bookkeeping edge cases, but I trust that anyone dedicated enough to include those in their formalism will be more than capable of resolving those issues on their own. Normally the task of setting out the bookkeeping for many-sorted structures is similarly relegated, but for our purposes here it will be prudent to consider it immediately.

For computable metric signatures, obviously we should require that the predicate ranges and maximum sort diameters be uniformly computable before recasting in the form above (although really all we need are uniformly computable upper and lower bounds), in order to ensure that we can uniformly compute presentations of recast structures from presentations of the original structures.

Definition 2.2. (i) A *metric signature* $\mathcal{L}$, is a tuple $(\mathcal{S}, \mathcal{P}, a, \Delta)$, where

- $\mathcal{S}$ is a set of sort symbols;
- $\mathcal{P}$ is a set of predicate symbols;
- $a : \mathcal{P} \to \mathcal{S}^{<\omega}$ is the arity function that assigns to each predicate symbol its finite string of input sorts (by an abuse of notation we will use a for formulas as well as atomic predicates); and
- for each predicate symbol P, $\Delta_P : [0,1] \to [0,1]$ is the syntactic modulus of uniform continuity of P.

(ii) A *computable metric signature* is a metric signature such that $\mathcal{S}$ and $\mathcal{P}$ are computable subsets of ω, a is a computable function which is total on $\mathcal{P}$, and $P \mapsto \Delta_P$ is a uniformly computable family of total computable functions.

Although in full generality moduli of uniform continuity can be specified as functions of each variable individually, on the level of a metric signature nothing is gained by such a generalization. Likewise there is no particular reason for moduli of uniform continuity to be continuous anywhere other than 0, but again nothing is gained and continuity is a more natural convention in the context of computable metric signatures.

The phrase 'syntactic modulus of uniform continuity' refers to the fact that in a given $\mathcal{L}$-structure the corresponding predicate may obey a stricter modulus of uniform continuity.

The definitions of restricted $\mathcal{L}$-formulas, $\mathcal{L}$-structures, and other such things is given in [5]. We should be clear about what a computable metric structure is.

Definition 2.3. Given a computable metric signature $\mathcal{L}$, a *computable $\mathcal{L}$-structure* is an $\mathcal{L}$-structure whose universes are a uniformly computable family of computable metric spaces (in the sense of [8]) and whose predicate interpretations are all uniformly computable functions.

Finally we will need a syntactically uniform notion of definable predicate, similar to the one given in [2].

Definition 2.4. (i) For a metric signature $\mathcal{L}$, a *finitary $\mathcal{L}$-formula* is an expression of the form $\sum_{n<\omega} 2^{-(n+1)}\varphi_n$, with φ_n a sequence of $[0,1]$-valued restricted $\mathcal{L}$-formulas such that the entire sequence contains finitely many free variables. Such a formula has a syntactic modulus of uniform continuity of $\sum_{n<\omega} 2^{-(n+1)}\Delta_{\varphi_n}$.

(ii) An *ω-infinitary $\mathcal{L}$-formula* is an expression of the same form without the restriction on involving finitely many free variables. (Such an expression has a uniformly computable syntactic modulus of uniform continuity in terms of the appropriate metric on ω-tuples, but it is somewhat more complicated to state.)

(iii) An *$\mathcal{L}$-formula* is either a finitary or an ω-infinitary $\mathcal{L}$-formula.

(iv) A *computable $\mathcal{L}$-formula* is an $\mathcal{L}$-formula such that the sequence of formulas φ_n is computable (the φ_n are required to be restricted formulas and can therefore be encoded by natural numbers).

Normally, definable predicates are defined relative to a single metric structure or an elementary class of metric structures in terms of uniformly convergent limits of restricted formulas, but the notion of formula given here is purely syntactic and can be interpreted in any $\mathcal{L}$-structure.

With some straightforward work, one can check that any definable predicate in the typical sense can be written in this form and that the resulting family of formulas is closed under this 'infinitary connective,' as well as all ordinary connectives, up to logical equivalence. In particular, even though continuous logic has an infinitary conjunction, it does not have a tall hierarchy of infinitary formulas the same way, for instance, $\mathcal{L}_{\omega_1\omega}$ does.

3 Expansions

We need to specify a few notions of expansions and interdefinability in continuous logic.

Definition 3.1. (i) For a given metric signature $\mathcal{L}$ and a finitary $\mathcal{L}$-formula $\varphi(\overline{x})$, a *definitional expansion* of $\mathcal{L}$ by φ is a metric signature $\mathcal{L}^*$ containing the same sorts as $\mathcal{L}$ and a single new predicate symbol P with $a(\varphi) = a(P)$ and $\Delta_\varphi = \Delta_P$. For $\mathfrak{A}$ an $\mathcal{L}$-structure, the corresponding $\mathcal{L}^*$-structure $\mathfrak{A}^*$ is given by interpreting P as φ. We also refer to iterated definitional expansions as definitional expansions.

(ii) An $\mathcal{L}$-structure, $\mathfrak{A}$, and a $\mathcal{K}$-structure, $\mathfrak{B}$, are *interdefinable* if there are definitional expansions $\mathfrak{A}^*$ and $\mathfrak{B}^*$ which make them isomorphic up to relabeling of sorts and predicate symbols (we allow metrics to be relabeled). A class C_0 of $\mathcal{L}$-structures and a class C_1 of $\mathcal{K}$-structures are interdefinable if there is a bijection between their isomorphism classes for which each pair is uniformly interdefinable via some fixed definitional expansions of signatures $\mathcal{L} \to \mathcal{L}^*$ and $\mathcal{K} \to \mathcal{K}^*$ and relabeling. (Note that we aren't requiring that the syntactic moduli of continuity match.)

(iii) Given a metric structure $\mathfrak{A}$, an *imaginary expansion* of $\mathfrak{A}$ is one of the following operations:

- Appending a product sort $P = \prod_{i<n} O_i$ for some $0 < n \leq \omega$ and sorts $O_i \in \mathcal{S}$. By convention the metric on a finitary product sort will always be the maximum of the component metrics and the metric on an ω-product sort will always be $\sup_{i<\omega} 2^{-i} d_{O_i}$. We also append projection predicates π_i on $P \times O_i$ for each $i < n$, where $\pi_i(\langle x_0, x_1, \ldots, x_{n-1}\rangle, y) = d_{O_i}(x_i, y)$.
- Appending a $\varnothing$-definable set D in sort O as a new sort O_D as well as an inclusion predicate ι on $O_D \times O$, where $\iota(x, y) = d_O(z, y)$ for $x \in O_D$ and $y, z \in O$ with z the element of D corresponding to x. The metric d_{O_D} is the restriction of d_O to D.
- For ρ, a $\varnothing$-definable pseudo-metric on sort O, appending the quotient sort O/ρ along with a quotient predicate q on $O \times O/\rho$, where $q(x, [y]_\rho) = \rho(x, y)$ for $x, y \in O$, where $[y]_\rho$ is the ρ-equivalence class of y. This is well-defined because ρ is a pseudo-metric.

Recall that we have restricted ourselves to relational languages at this point, which is why the projection, inclusion, and quotient maps are encoded as predicates.

We also refer to iterated imaginary expansions as imaginary expansions. [1]

The added generality of allowing ω-tuples and passing to definable sets is natural and somewhat necessary in continuous logic[5, Section 11]. ω-tuples are necessary for canonical parameters since a formula can involve countably many parameters, but note that for any formula φ on an ω-product sort, if $\mathfrak{A} \models \varphi(\overline{a})$, then for any $\varepsilon > 0$ the fact that $\mathfrak{A} \models \varphi(\overline{a}) < \varepsilon$ only depends on finitely many terms in $\overline{a}$, uniformly as a function of ε, because φ needs to be uniformly continuous with regards to the ω-product metric. Because of this, ω-product sorts are just as safe as finitary product sorts in terms of compatibility with ultraproducts and preserving the category of models. Explicitly passing to definable sets is necessary in situations such as the following: In a connected metric structure $\mathfrak{M}$ with a non-trivial definable discrete subset D, there is no uniformly continuous pseudo-metric ρ on $\mathfrak{M}$ that will make $\mathfrak{M}/\rho$ isometric to D (or D plus a single new point or anything else you would do in the discrete setting), since the quotient map $\mathfrak{M} \to \mathfrak{M}/\rho$ is continuous and continuous functions preserve connectedness.

Lemma 3.2. *(i) For any metric signature $\mathcal{L}$ (not necessarily countable), there is a metric signature $\mathcal{K}$ which is interdefinable with an imaginary expansion of $\mathcal{L}$ such that $\mathcal{K}$ has a uniform bound of 2 on the arities of its predicate symbols. For computable signatures, the signature $\mathcal{K}$ is uniformly computable from $\mathcal{L}$ and presentations of $\mathcal{L}$-structures can be uniformly converted into corresponding presentations of $\mathcal{K}$-structure and vice versa.*

[1] Even though imaginary expansions are defined for structures and not signatures, of the three forms of imaginary expansion, only expansion by a definable set is not uniform across all structures of a given signature, as every definable pseudo-metric can be written in the form $\rho(x, y) = \sup_{\overline{z}} |\varphi(x, \overline{z}) - \varphi(y, \overline{z})|$, as $\rho(x, y) = \sup_z |\rho(x, z) - \rho(y, z)|$, and such an expression is a definable pseudo-metric in any $\mathcal{L}$-structure.

(ii) There is a $\mathcal{K}$-theory $T_{\mathcal{L}}$, uniformly computable from $\mathcal{L}$, such that the models of $T_{\mathcal{L}}$ are precisely the interpretations of $\mathcal{L}$-structures.

Proof. (i) For each predicate symbol p, we can define a unary formula on the sort $\Pi_{O\in a(p)}O$ in the obvious way. These, together with projection maps between product sorts and the original $\mathcal{L}$-sorts, are clearly enough to define any predicate originally definable in an $\mathcal{L}$-structure in a completely uniform way. Since the projection maps are encoded as 2-ary predicates, we have the required arity bound. This procedure is also clearly uniformly computable, both for signatures and presentations of structures.

(ii) All that $T_{\mathcal{L}}$ needs to say is that the predicates corresponding to projection maps are actually projection maps and that the product sorts are products of the sorts they project onto. $\square$

Definition 3.3. If $\mathcal{L}$ is a metric signature with designated home sort H and $0 < r \leq 1$ is a real number, $C_{\mathcal{L},r}$ is the class of $\mathcal{L}$-structures $\mathfrak{A}$ satisfying $\mathrm{diam}(H^{\mathfrak{A}}) \geq r$.

The following lemma is the source of all non-uniformity relative to r in the entire construction and is analogous to the fact that a discrete structure with only one element cannot interpret any structure with more than one element. It could be avoided by appending a new compact sort isometric to $[0,1]$ with the standard metric, or literally any other fixed non-trivial compact metric space, and letting that be the designated home sort H.

Lemma 3.4. *Let X be a compact metric space. There is a uniformly definable imaginary Y such that for any $\mathfrak{A} \in C_{\mathcal{L},r}$, $Y^{\mathfrak{A}} \cong X$, with each point of $Y^{\mathfrak{A}}$ and every continuous function $(Y^{\mathfrak{A}})^n \to [0,1]$ uniformly $\varnothing$-definable.*

Proof. Let x_0, x_1, y_0, y_1 be variables in H and consider the $\mathcal{L}$-formula

$$\rho(x_0, x_1, y_0, y_1) = \frac{1}{r}|d(x_0, x_1) - d(y_0, y_1)| \wedge 1.$$

This is a pseudo-metric on H^2. H^2/ρ contains more than one point for any $\mathfrak{A} \in C_r$, because of the diameter requirement, in particular it has a definable subset consisting of the ρ-equivalence classes of pairs satisfying $d(x_0, x_1) = 0$ and pairs satisfying $d(x_0, x_1) \geq r$, with each of those points being $\varnothing$-definable by the formulas $\frac{1}{r}d(x_0, x_1)$ and $1 \dot{-} \frac{1}{r}d(x_0, x_1)$, respectively. Let D denote this definable set. Clearly D is always isometric to the discrete space with two points so in particular $C = D^{\omega}$ is an isometric copy of Cantor space with the standard metric with every point uniformly $\varnothing$-definable. It is well known that Cantor space continuously surjects onto any compact metric space (X, d^X), so by pulling back d^X to C^2 we get a continuous pseudo-metric on C whose quotient is isometric to X. Therefore, since the type space $S_{C^2}(T)$ is isomorphic to C^2 (both metrically and topologically), the pullback metric is a continuous function on $S_{C^2}(T)$ and is thus a definable pseudo-metric on C. Since each point of C is uniformly definable, this gives the required uniformly definable imaginary Y.

Finally for an arbitrary continuous function $f : X^n \to [0,1]$, the pullback on the type space $S_{C^n}(T)$ is continuous and therefore definable. By construction it is compatible with the quotient map $C \to Y$ and is therefore a definable predicate on the imaginary $Y^{\mathfrak{A}}$.

$\square$

There are some potential subtleties involving uniform computability of formulas definining computable compact imaginaries and computable predicates on them. In the current context we only need Lemma 3.4 for a small handful of very specific tame compact metric spaces, so we'll deal with computability on a case-by-case basis.

Lemma 3.5. *For any $C_{\mathcal{L},r}$, with $r > 0$, if $\{O_n^{\mathfrak{A}}\}_{n<k}$ is a finite collection of sorts of diameter ≤ 1, then the disjoint union $U = \bigsqcup_{n<k} O_n$ with metric $d(x,y) = 1$ for $x \in O_n$ and $y \in O_m$ with $n \neq m$ and $d(x,y) = d_{O_n}(x,y)$ for $x, y \in O_n$ is a uniformly definable imaginary in $C_{\mathcal{L},r}$. Furthermore the formulas defining U are uniformly computable in $\mathcal{L}$, r, and the list of sorts.*

Proof. By Lemma 3.4, the discrete space $\Delta_k = \{0, ..., k-1\}$, with the metric $\delta(x,y) = 1$ if $x \neq y$, is uniformly an imaginary of $C_{\mathcal{L},r}$ (although in particular we don't have to go through Cantor space and we can realize Δ_k as a quotient of some $\Delta_{2^\ell} = (\Delta_2)^\ell$ in a clearly uniformly computable way). Furthermore we can arrange that each element of Δ_k is definable.

Define a formula $\rho(\overline{x}, \overline{y})$ on $\Delta_k \times \prod_{n<k} O_n$ by

$$\rho(\overline{x}, \overline{y}) = \left(\delta(x_0, y_0) + \bigwedge_{n<k} \left(d(x_{n+1}, y_{n+1}) + \delta(x_0, n) \right) \right) \wedge 1.$$

Checking definitions gives that $\Delta_k \times \prod_{n<k} O_n / \rho$ is the required imaginary. This formula is also clearly uniformly computable. $\square$

4 Countable disjoint unions of sorts

A common trick in discrete logic is merging a finite collection of sorts by taking the disjoint union and adding unary predicates selecting out each sort. This can't be extended to infinitely many predicates without changing the category of models: by compactness there will be models with elements not in any given sort. The added flexibility of continuous logic allows us to do this with countably many sorts at once without changing the category of models, specifically we can arrange it so that any sequence of types that ought to limit to an 'unsorted' type is shunted into a single unique overflow point. This is very similar to the 'enboundment' method used in [3] to treat unbounded metric structures.

It should be noted that if $\mathcal{L}$ has finitely many sorts and (possibly infinitely many) predicates with uniformly bounded arity, this section can be skipped and the construction in Theorem 6.2 will work directly.

Definition 4.1. Let $\{O_n\}_{n<\omega}$ be a countable sequence of $\mathcal{L}$-sorts. For any $\mathcal{L}$-structure $\mathfrak{A}$, the *countable metric disjoint union* of $\{O_n\}_{n<\omega}$ is a metric structure with the following set as its universe:

$$U^{\mathfrak{A}} = \{*\} \cup \bigsqcup_{n<\omega} O_n^{\mathfrak{A}}$$

where $*$ is a single new point. The metric disjoint union is written $\bigsqcup^*_{n<\omega} O_n$. $U^{\mathfrak{A}}$ has the following metric: let $x, y \in O_n^{\mathfrak{A}}$, $z \in O_m^{\mathfrak{A}}$, with $n \neq m$, then

$$d_U^{\mathfrak{A}}(x,y) = 2^{-n} d_{O_n}^{\mathfrak{A}}(x,y),$$

$$d_U^{\mathfrak{A}}(x,z) = |2^{-n} - 2^{-m}|,$$

$$d_U^{\mathfrak{A}}(x,*) = 2^{-n},$$

where the other values are determined by symmetry. We will prove in Proposition 4.2 that this defines a complete metric space.

A predicate on some $O_{n_1}^{\mathfrak{A}} \times \cdots \times O_{n_k}^{\mathfrak{A}}$ is extended to a predicate on $U^{\mathfrak{A}}$ by setting its value to 1 (i.e. 'false') when the input is not part of its domain.

Finally we add a distance predicate for the set $\{*\}$ (recall that we have restricted ourselves to relational languages, so we can't use a constant).

Proposition 4.2. *(i) The countable metric disjoint union, $U = \bigsqcup^*_{n<\omega} O_n$, of a sequence $\{O_n\}_{n<\omega}$ of $\mathcal{L}$-sorts is well-defined, i.e. the metric given in the definition is actually a metric and defines a complete metric space.*

(ii) The predicates interpreted on it are uniformly continuous. In particular if they are Lipschitz in the original signature they will still be Lipschitz (although possibly with a different Lipschitz constant).

(iii) The countable metric disjoint union is isomorphic to a uniformly definable imaginary for all $\mathfrak{A} \in C_{\mathcal{L},r}$. The relevant formulas and the map of presentations $\mathfrak{A} \mapsto U^{\mathfrak{A}}$ are uniformly computable from the sequence $\{O_n\}_{n<\omega}$, the signature $\mathcal{L}$, and the real number r, so in particular if all of those are computable, then the relevant formulas and the map of presentations are computable.

(iv) Each O_n as a subset of U is a definable subset of U and (considering U as an imaginary sort) there is a definable bijection between O_n as a sort and O_n as a definable subset of U. The relevant formulas are uniformly definable in $\mathcal{L}$ and computable.

(v) For a fixed sequence $\mathscr{S} = \{O_n\}_{n<\omega}$ of $\mathcal{L}$-sorts with $O_0 = H$, the designated home sort, there is a signature $\mathcal{L}_{\mathscr{S}}$ and a theory $T_{\mathscr{S}}$, uniformly computable from $\mathcal{L}$ and $\mathscr{S}$, and a $\mathcal{L}_{\mathscr{S}}$-sentence $\Xi_{\mathscr{S}}$, such that the models of $T_{\mathscr{S}} \cup \{r \dot{-} \Xi_{\mathscr{S}}\}$ are precisely the same as reducts to the sort $\bigsqcup^*_{n<\omega} O_n$ of structures in $C_{\mathcal{L},r}$

Proof. (i) The expression given for d clearly obeys all metric space axioms besides the triangle inequality. The only unobvious case is the one consisting of two points in some O_n and a third point in some O_m with $n \neq m$, and letting $O_\omega = \{*\}$ with the understanding that "$2^{-\omega}$" $= 0$. Let $x, y \in O_n$ and $z \in O_m$ with $n \neq m$. By symmetry there are only 2 cases to check:

- $d(x,y) \leq 2^{-n}$ and $d(x,z) = d(y,z) = |2^{-n} - 2^{-m}| \geq 2^{-(n \wedge m+1)}$, so $d(x,z) + d(y,z) \geq 2^{-(n\wedge m)} \geq 2^{-n} \geq d(x,y)$ and in this case the triangle inequality is obeyed.
- $d(x,z) = |2^{-n} - 2^{-m}|$, so $d(x,z) \leq d(x,y) + |2^{-n} - 2^{-m}| = d(x,y) + d(y,z)$.

So the triangle inequality is obeyed in all cases.

To see that the metric space is complete, note that any Cauchy sequence is either eventually contained in some O_n or is limiting to $*$.

(ii) If a predicate P on sort $O_{n_1} \times \cdots \times O_{n_k}$ has modulus of uniform continuity $\Delta_P(x)$, then the corresponding predicate on U is uniformly continuous with modulus of uniform continuity

$$\Delta_P^*(x) = \big(\Delta_P((2^N x) \wedge 1) \vee (2^{N+1}x)\big) \wedge 1,$$

where $N = n_0 \vee \cdots \vee n_{k-1}$. Note that if P has Lipschitz constant L, then on U it will have Lipschitz constant $2^{N+1}L$, and in particular it will still be Lipschitz.

(iii) By Lemma 3.4, the class $C_{\mathcal{L},r}$ has a uniformly definable imaginary isometric to the metric space (X,d) where $X = \{0\} \cup \{2^{-n} : n < \omega\}$ and d is the standard metric on $\mathbb{R}$. Let $W = X \times \Pi_{n<\omega} O_n$ be the infinitary product sort.

(iv) For any $x \in U$, $d(x, O_n) = |d(x,*) - 2^{-n}|$.

Let $Q : X \to [0,1]$ be the natural inclusion map, which is a definable predicate on X uniformly for all members of $C_{\mathcal{L},r}$. For each n, let

$$R_n(x) = 1 \dot{-} 2^{n+1}|Q(x) - 2^{-n}|,$$

i.e. R_n is a predicate on X which takes on the value 1 at 2^{-n} and 0 everywhere else. Now define a pseudo-metric on W by

$$\rho(\overline{x}, \overline{y}) = |Q(x_0) - Q(y_0)| + \sum_{n<\omega} 2^{-n} R_n(x_0) R_n(y_0) d_{O_n}(x_{n+1}, y_{n+1}).$$

Although in principle this is $[0,2]$-valued, by construction it will only take on values in $[0,1]$. Taking the quotient W/ρ will identify any two elements $\overline{a}, \overline{b} \in W$ if and only if $a_0 = b_0$ and either $a_0 = 0$ or $a_0 = 2^{-n}$ and $a_n = b_n$. So by making the identification of elements of the form $(2^{-n}, \dots, a_n, \dots)$ with a_n and elements of the form $(0, \dots)$ with $*$, we get a bijection between W/ρ and U, and by checking definitions we see that ρ induces the correct metric on U.

(v) The signature $\mathcal{L}_{\mathscr{S}}$ has a single sort the same predicate symbols as $\mathcal{L}$ with the same total arity along with a single new unary predicate symbol Q. For each predicate symbol P, the syntactic modulus of continuity is $\big(\Delta_P((2^N x) \wedge 1) \vee (2^{N+1}x)\big) \wedge 1$, where Δ_P is the syntactic modulus of continuity of P in $\mathcal{L}$, and $\Delta_Q(x) = x$.

Recall that for a given formula φ there are axioms determined by φ that hold exactly when φ is the distance predicate of a definable set[5, Theorem 9.12]. Since these axioms are fixed restricted sentences depending on φ, they are uniformly computable from φ. $T_{\mathscr{S}}$ has axioms saying that Q is a distance predicate and that the set defined by Q is non-empty and has diameter 0. Let

$*$ be a constant referring to the unique point defined by Q in order to make the following axioms easier to write down.

Let $f : [0,1] \to [0,1]$ be a total computable continuous function whose zeroset is precisely $X = \{0\} \cup \{2^{-n} : n < \omega\}$. $T_{\mathscr{S}}$ has the axioms

$$\sup_x f(d(x,*)),$$

$$\inf_x |d(x,*) - r|, \text{ for each } r \in X.$$

For each $r \in X$, $T_{\mathscr{S}}$ has axioms stating that $|d(x,*) - r|$ is a distance predicate for a definable set, by abuse of notation label those definable sets O_n. The sentence $\Xi_{\mathscr{S}}$ is given by

$$\Xi_{\mathscr{S}} = \sup_{x,y \in O_0} d(x,y),$$

i.e. the diameter of O_0. Finally there are axioms for each predicate symbol P and incorrect sequence of input sorts saying that P takes on the value 1 on those inputs (this is expressible since the sorts are definable sets) and axioms stating that P obeys the modulus of continuity Δ_P relative to the rescaled metrics on the O_n. $\square$

The following proposition is clear by construction and in particular part (iii) of Proposition 4.2 above.

Proposition 4.3. *If $\mathcal{L}$ is a metric signature with countably many sorts and we let $U = \bigsqcup^*_{O \in \mathcal{S}} O$ be the imaginary disjoint union of all $\mathcal{L}$-sorts, then for all $\mathfrak{A} \in C_{\mathcal{L},r}$, $\mathfrak{A}$ and $U^{\mathfrak{A}}$ have uniformly definable imaginary expansions which are uniformly interdefinable.*

Aside from the issue of topological dimension and continuous degrees of points in the structure discussed in the introduction, one of the mild technical advantages of a countable metric disjoint union over an ω-product is that parameters in non-trivial ω-products tend to be poorly behaved in that they act like countable collections of parameters rather than finite collections of parameters. This general phenomenon of single parameters acting like countable collections of parameters can be blamed for many of the pathologies in continuous logic (e.g. pairs ab such that $\mathrm{tp}(ab)$ is principal but $\mathrm{tp}(a/b)$ is not, theories with exactly two separable models, small theories with only 'approximately ω-saturated' separable models, and ω-categorical theories which fail to be ω-categorical after naming an element). In [1], Usvyatsov and Ben Yaacov introduced the notion of a d-finite type, which intuitively speaking characterizes when a finitary type actually behaves like a discrete finitary type, rather than a discrete ω-type. Uniform d-finiteness is a technical strengthening of d-finiteness that was needed in an analog of Lachlan's theorem on the number of countable models of a superstable theory in [1].

Proposition 4.4. *(i) Let $\overline{a} \in \bigsqcup^*_{n<\omega} O_n$ be an ℓ-tuple of elements not equal to $*$. For any set B of parameters, $tp(\overline{a}/B)$ is (uniformly) d-finite as a type in the correct product sort if and only if it is (uniformly) d-finite as a type in the sort $\left(\bigsqcup^*_{n<\omega} O_n\right)^\ell$. (Note that since $* \in dcl(\varnothing)$, its type is always uniformly d-finite and adding it to a tuple preseves d-finiteness.)*

(ii) For any (locally) compact set $B \subset O^{\mathfrak{A}}_k$, the corresponding set in $\bigsqcup^*_{n<\omega} O^{\mathfrak{A}}_n$ is (locally) compact. (Although note that the countable metric disjoint union will typically fail to be locally compact at $*$.)

(iii) For any topologically finite dimensional (resp. weakly infinite dimensional) set $B \subset O^{\mathfrak{A}}_k$, the corresponding set in $\bigsqcup^*_{n<\omega} O^{\mathfrak{A}}_n$ is finite dimensional (resp. weakly infinite dimensional). If each $O^{\mathfrak{A}}_k$ is finite dimensional, then $\bigsqcup^*_{n<\omega} O^{\mathfrak{A}}_n$ will be either finite dimensional or weakly infinite dimensional and locally finite dimensional away from $*$. If each $O^{\mathfrak{A}}_k$ is weakly infinite dimensional, then $\bigsqcup^*_{n<\omega} O^{\mathfrak{A}}_n$ is as well.

Proof. These all follow from the fact that the natural inclusion maps $O_k \to \bigsqcup^*_{n<\omega} O_n$ are open, isometric-up-to-scaling, and bijections between definable sets. □

In particular if T is 'hereditarily ω-categorical' (i.e. ω-categorical over every finite set of parameters) or has an exactly ω-saturated separable model, then $\mathrm{Th}\left(\bigsqcup^*_{n<\omega} O_n\right)$ will as well[1].

5 Making everything Lipschitz

Ultimately we will need all of our predicate symbols to be Lipschitz since they will be encoded directly into a metric and metrics are always Lipschitz. There are a couple of ways to accomplish this. If the reader does not care about computability, this section can be skipped using the following fact. Also it should be noted that Fact 5.1 works for uncountable metric signatures, but the result that we will use, Proposition 5.6, does not in general.

Fact 5.1. *Let (X, d) be a metric space and $f : X \to [0, 1]$ be a uniformly continuous function. For each $0 < n < \omega$, let*

$$f_n(x) = \inf_y \left(\frac{1}{n} f(y) + d(x, y)\right) \wedge 1.$$

Then $f_n(x)$ is a sequence of 1-Lipschitz functions such that $nf_n \to f$ uniformly as $n \to \infty$.

In general that procedure would cost a jump to compute on a given structure, i.e. if some degree $\mathbf{a}$ computes a structure (M, P) with predicate P, then $\mathbf{a}'$ will compute $(M, P_0, P_1, \dots)$ with P_n given by the formula in Fact 5.1, so to ensure that the construction is computable, we will have to use something else. The idea is that if α is a concave non-decreasing function such that $\alpha(0) = 0$, then for any metric d, $\alpha(d)$ is also a metric. So if one of our predicates P has a concave non-decreasing modulus of uniform continuity, then we can compose it with the metric to get a uniformly equivalent metric relative to which P is 1-Lipschitz.

The following is a fairly elementary real analytic fact, but we will include a proof for the sake of demonstrating that the procedure is computable. Note that we could avoid this lemma entirely if we require that moduli of continuity be non-decreasing and sub-additive, which is often required for moduli of uniform continuity and can always be arranged as per this lemma.

Lemma 5.2. *Let $\delta : [0,1] \to [0,1]$ be a continuous function satisfying $\delta(0) = 0$. There is a continuous, concave, non-decreasing function $\alpha : [0,1] \to [0,1]$ satisfying $\alpha(0) = 0$ and $\alpha \geq \delta$. Furthermore, α is uniformly computable from δ.*

Proof. α will be the 'non-decreasing convex hull of δ,' defined by the following formula:

$$\alpha(x) = \inf\{mx + b \,:\, 0 \leq m, b,\, (\forall y \in [0,1]) my + b \geq \delta(y)\}.$$

For each $n < \omega$, define

$$\alpha_n(x) = \inf\{mx + b \,:\, 0 \leq m, b,\, (\forall k \in \{0, 1, \dots, 2^n\}) m(2^{-n}k) + b \geq \delta(2^{-n}k)\}.$$

When computing α_n, the largest m necessary is at most

$$m_n = 2^n \sup_{0 \leq k < 2^n} |\delta(2^{-n}(k+1)) - \delta(2^{-n}k)|,$$

and the largest b is always at most 1, so the computation of α_n amounts to minimizing a δ-computable linear function on a δ-computable bounded polytope, so the α_n are uniformly computable in δ[8, Chapter 5]. Furthermore note that since each α_n is the infimum of a family of Lipschitz functions with uniformly bounded Lipschitz coefficients, α_n is Lipschitz and in particular continuous.

Now all we need to show is that α_n converges uniformly to α with a computable modulus of uniform convergence. For computability considerations, we will need the fact that the modulus of uniform continuity of a continuous function f on $[0,1]$ is uniformly computable from f [8, Chapter 6]. Let Δ_δ be the modulus of uniform continuity of δ. By replacing Δ_δ with $\sup_{0 \leq y \leq x} \Delta_\delta(y)$ (with is uniformly computable from Δ_δ, since $[0,x]$ is effectively compact uniformly in x), we may assume that Δ_δ is non-decreasing.

Now note that for each $n < \omega$, we have the following inequality:

$$\alpha_n \leq \alpha \leq \alpha_n + 2\Delta_\delta(2^{-n}). \qquad (\star)$$

To see that this inequality is true, note that for each interval $I = [2^{-n}k, 2^{-n}(k+1)]$, we must have

$$\delta(x) \leq \left[\delta(2^{-n}k) \vee \delta(2^{-n}(k+1))\right] + \Delta(2^{-n}) \leq \left[\delta(2^{-n}k) \wedge \delta(2^{-n}(k+1))\right] + 2\Delta(2^{-n}),$$

for all $x \in I$. Therefore, if $m, b \geq 0$ satisfy the requirements in the infimum defining α_n, then for all $x \in I$,

$$mx + b \geq \delta(2^{-n}k) \vee \delta(2^{-n}(k+1)),$$

and thus ($\star$) follows, so we get that $\alpha_n \to \alpha$ uniformly as $n \to \infty$, so α is continuous. Furthermore, we clearly have a uniformly computable modulus of uniform convergence, so α is uniformly computable.

Finally note that α is concave and non-decreasing by construction (these are preserved by infima) and $\alpha(0) = 0$ since for every $\varepsilon > 0$, there is an $m > 0$ such that $mx + \varepsilon \geq \delta(x)$ for all $x \in [0, 1]$, by continuity of δ. □

So if we have a single modulus of continuity that works for all relation symbols we can find an inter-definable structure with a Lipschitz signature. Fortunately we can always arrange this if our signature is countable.

Definition 5.3 (Uniform uniform continuity)**.** (i) A family of functions $f \in F$ on a metric space X is *uniformly uniformly continuous* if there is a single modulus of uniform continuity valid for all $f \in F$.

(ii) A metric signature $\mathcal{L}$ is *uniformly uniformly continuous* if $\Delta_P = \Delta_Q$ for all predicate symbols P and Q.

Recall that two metric spaces (X_0, d_0), (X_1, d_1) are bi-uniformly isomorphic if there is a uniformly continuous bijection $f : X_0 \to X_1$ with uniformly continuous inverse. Two metrics d_0, d_1 on the same space X are uniformly equivalent if (X, d_0) and (X, d_1) are bi-uniformly isomorphic under the identity map.

Lemma 5.4. *(i) If d is a $[0,1]$-valued metric and $\alpha : [0, 1] \to [0, 1]$ is a continuous, concave, non-decreasing function satisfying $\alpha(0) = 0$, then $\alpha(d) \vee d$ is a metric that is uniformly equivalent to d.*

(ii) If (X, d) is a metric space with diameter ≤ 1 and $f_i : X \to [0, 1]$, $i \in I$, is a family of uniformly uniformly continuous functions with continuous, sub-additive, non-decreasing modulus of uniform continuity α, then $(X, \alpha(d) \vee d)$ is a metric space bi-uniformly isomorphic to (X, d), such that the family $\{f_i\}_{i \in I}$ is 1-Lipschitz.

Proof. (i) Concave functions are sub-additive. All of the pseudo-metric axioms are preserved under composition with sub-additive, non-decreasing functions which fix 0, so $\alpha(d)$ is a pseudo-metric. The maximum of two pseudo-metrics is still a pseudo-metric, so $\alpha(d) \vee d$ is a pseduo-metric. $\alpha(d) \vee d = 0$ if and only if $d = 0$, so it is actually a metric. $\alpha(d) \vee d$ and d are clearly uniformly equivalent.

(ii) This is immediate from (i). □

In the previous lemma we only need to take the maximum with d on the off chance that $\alpha = 0$ and there's no harm in doing so.

Lemma 5.5. *If $\mathcal{L}$ is a countable metric signature, then it is interdefinable with a uniformly uniformly continuous metric signature $\mathcal{K}$. Furthermore if $\mathcal{L}$ is computable, then we can take $\mathcal{K}$ to be uniformly computable in $\mathcal{L}$.*

Proof. Let $\{P_i\}_{i<\omega} = \mathcal{P}$ be an enumeration of all the predicate symbols in $\mathcal{L}$ (in any sort). For each $i < \omega$, let Q_i be the $\mathcal{L}$-formula $2^{-(i+1)}P_i$. The $\mathcal{L}$-formulas Q_i are uniformly uniformly continuous with regards to the modulus of uniform continuity $\Delta = \sum_{i<\omega} 2^{-(i+1)}\Delta_{P_i}$. If we let $\mathcal{K}$ be a metric signature with the same sorts as $\mathcal{L}$ and predicate symbols for the Q_i, each with $\Delta_{Q_i} = \Delta$, then $\mathcal{K}$ is the required metric signature.

The procedure described in Lemma 5.2 is uniformly computable, so passing from $\mathcal{L}$ to $\mathcal{K}$ is uniformly computable as well. $\square$

Proposition 5.6. *(i) If $\mathcal{L}$ is a countable metric signature, then it is interdefinable with a* 1*-Lipschitz metric signature $\mathcal{K}$, i.e. a signature such that $\Delta_P(x) = x$ for all predicate symbols P (although not for metrics, which are necessarily* 2*-Lipschitz). Furthermore $\mathcal{K}$ is uniformly computable from $\mathcal{L}$.*

(ii) There is a $\mathcal{K}$ theory $T_{\mathcal{L}}$ such that the models of $T_{\mathcal{L}}$ are precisely the interpretations of $\mathcal{L}$-structures as $\mathcal{K}$-structures. Furthermore $T_{\mathcal{L}}$ is uniformly computable from $\mathcal{L}$.

Proof. (i) Aside from what we have already outlined in this section, the only subtlety is that the passage from d to $\alpha(d) \vee d$ may delete some information contained in d due to 'clipping' wherever α is locally constant. To remedy this all we need to do is add a new binary 1-Lipschitz predicate symbol $P_{d,O}$ for each sort O whose interpretation is $\frac{1}{2}d_O$ before running the construction in this section. This doesn't prevent α from clipping the metric, but we lose no information since we have an unmodified copy of the original metric as a predicate.

(ii) $T_{\mathcal{L}}$ just needs to express that every predicate symbol is uniformly continuous with regards to the original metrics $d_O = 2P_{d,O}$ in the appropriate way, i.e. with axioms of the form

$$\sup_{x,y} |P(x) - P(y)| \dot{-} \Delta_P(2P_{d,O}(x,y)),$$

and analogous axioms for predicates on more than one sort. $\square$

6 Encoding in metric spaces

Most of the coding tricks used in the two following constructions boil down to the fact that if X and Y are metric spaces with diameter ≤ 1, then for any 1-Lipschitz function $f : X \times Y \to [0,1]$, you can put a metric on $X \sqcup Y$ by setting $d(x,y) = 2 + f(x,y)$ for $x \in X$ and $y \in Y$, so assuming we can define X and Y, we can recover f from the metric alone. The other fundamentally important thing is that since our metric structures have bounded diameter, we can add points at a larger diameter to ensure that they are $\varnothing$-definable regardless of the content of the embedded metric structure.

Theorem 6.1. *(i) If $\mathcal{L}$ is a countable metric signature, then for any $C_{\mathcal{L},r}$, there is a uniformly definable imaginary X such that for any $\mathfrak{A} \in C_{\mathcal{L},r}$, $\mathfrak{A}$ and the* *purely metric reduct* $X_0^{\mathfrak{A}} = (X^{\mathfrak{A}}, d)$ *are uniformly bi-interpretable in the sense that*

- *there are uniformly definable imaginary expansions of* $\mathfrak{A}$ *and* $X_0^{\mathfrak{A}}$ *which are uniformly interdefinable, and*
- *there are uniformly definable bijections between the sorts of* $\mathfrak{A}$ *and definable subsets of* $X_0^{\mathfrak{A}}$*, and* $X_0^{\mathfrak{A}}$ *is contained in the definable closure of the images of those bijections.*

Furthermore the interpretation preserves embeddings and (uniform) d-finiteness of types. If the original structure is not strongly infinite dimensional, then the interpreted structure will also not be strongly infinite dimensional. The interpretation preserves local compactness and local finite dimensionality away from a fixed, compact, $\varnothing$*-definable set of bad points.*

(ii) For any countable metric signature $\mathcal{L}$*, there is a first-order theory* $T_{\mathcal{L}}$ *and a sentence* Ξ *such that for any* $r \in (0,1]$*, the class of metric spaces of the form* $X_0^{\mathfrak{A}}$ *for* $\mathfrak{A} \in C_{\mathcal{L},r}$ *is precisely the set of models of* $T_{\mathcal{L}} \cup \{r \dot{-} \Xi\}$*. If* $\mathcal{L}$ *is a computable signature, then* $T_{\mathcal{L}}$ *is computable.* Ξ *does not depend on* $\mathcal{L}$ *and is always computable.*

Furthermore there are computable mappings of presentations of $\mathcal{L}$*-structures to presentations of models of* $T_{\mathcal{L}}$ *and vice versa (these mappings do not depend on* r*).*

Proof. (i) By applying Lemma 3.2, we may assume that $\mathcal{L}$ has a uniform arity bound of 2. By applying Propositions 4.3 and 5.6, we may assume that $\mathcal{L}$ has a single sort and is 1-Lipschitz. By recasting unary predicates P as binary predicates using $P(x,y) = P(x)$, we may assume that all predicates are binary.

Let $\{P_n\}_{n<\omega}$ be an enumeration of all predicates with $P_0(x,y) = \frac{1}{2}d(x,y)$.

$X^{\mathfrak{A}}$ will have the set $A \sqcup A \times \omega \sqcup \{\infty, t\}$ as its universe, where $A \times \omega \sqcup \{\infty\}$ will be a modified countable metric disjoint union, with overflow point ∞, and t will be a tag to keep things straight. $X^{\mathfrak{A}}$ will have the unique metric defined by:

- $d(x,y) = d^{\mathfrak{A}}(x,y)$, for $x, y \in A$.
- $d(x,(y,n)) = 2 + 2^{-n-1}d^{\mathfrak{A}}(x,y)$, for $x \in A$ and $(y,n) \in A \times \omega$.
- $d(x,\infty) = 2$ for $x \in A$.
- $d(x,t) = 5$, for $x \in A$.
- $d((x,n),(y,n)) = 2^{-n}d^{\mathfrak{A}}(x,y)$, for $(x,n),(y,n) \in A \times \omega$.
- $d((x,n),(y,n+1)) = 2^{-n-1}(1 + P_n^{\mathfrak{A}}(x,y))$, for $(x,n),(y,n+1) \in A \times \omega$.
- $d((x,n),(y,m)) = |2^{-n} - 2^{-m}|$, for $(x,n),(y,m) \in A \times \omega$ with $|n-m| > 1$.
- $d((x,n),\infty) = 2^{-n}$

- $d((x,n),t) = 4 + 2^{-n-1}$, for $(x,n) \in A \times \omega$.
- $d(\infty, t) = 4$

All of the metric space axioms except for the triangle inequality are clearly obeyed by d. If all three points are in the same copy of A then the triangle inequality is obeyed, so we only need to check mixed triples. The majority of cases are mechanical to check, but there are a handful of tight or subtle cases that we will write out explicitly. Let $x, y \in A$ and $(z,n), (w,n), (u,n+1), (v,n+1), (s,m) \in A \times \omega$ with $|n-m| > 1$, where $n, m < \omega$. Also, recall that if $n \neq m$, then $|2^{-n} - 2^{-m}| \geq 2^{-n-1}$.

- $d(x,(z,n)) \leq 2 + 2 \cdot 2^{-n-2} \leq d(x,(u,n+1)) + d((u,n+1),(z,n))$
- $d((z,n),(u,n+1)) \leq 2^{-n-1}(1 + P_n(w,u) + d(z,w))$
  $\leq d((z,n),(w,n)) + d((w,n),(u,n+1))$
- $d((z,n),(w,n)) \leq 2 \cdot 2^{-n-2} \leq d((z,n),(u,n+1)) + d((u,n+1),(w,n))$
- $d((z,n),(v,n+1)) \leq 2^{-n-1}(1 + P_n(z,u) + d(u,v))$
  $= d((z,n),(u,n+1)) + d((u,n+1),(v,n+1))$
- $d((z,n),t) = 4 + 2 \cdot 2^{-n-2} \leq d((z,n),(u,n+1)) + d((u,n+1),t)$

Just as in the proof of Proposition 4.2, let $Y = 0 \cup \{2^{-n} : n < \omega\}$ and let $Q : Y \to [0,1]$ be the natural inclusion map, which is a definable predicate on Y. For each n, let

$$R_n(x) = 1 \dot{-} 2^{n+1}|Q(x) - 2^{-n}|,$$

and define a pseudo-metric, ρ, on $Y \times A$ by

$$\beta_n(\overline{x},\overline{y}) = R_n(x_0)R_{n+1}(y_0)P_n(x_{n+1},y_{n+2}) + R_{n+1}(x_0)R_n(y_0)P_n(y_{n+1},x_{n+2}),$$

$$\rho(\overline{x},\overline{y}) = |Q(x_0)-Q(y_0)| + \sum_{n<\omega} 2^{-n}\left(R_n(x_0)R_n(y_0)d(x_{n+1},y_{n+1}) + \frac{1}{2}\beta_n(\overline{x},\overline{y})\right).$$

Then $Y \times A/\rho$ will correspond to $A \times \omega \sqcup \{\infty\}$, where ∞ is the ρ-equivalence class of any element of the form $\langle 0, x\rangle$ for $x \in A$.

Recall that an element or set is definable if there is a formula which defines its distance predicate. If we have a $\{0,1\}$-valued indicator function, $\varphi(x)$, for the set $\varphi^{-1}(0)$, then that's even better and we can always define the distance to the set by $d(x,\varphi^{-1}(0)) = \inf_y d(x,y) + 5\varphi(y)$ if we need it. Once a point is definable we'll freely use it as a constant to make the following formulas simpler[5, Proposition 9.18]. We'll either find an indicator function or a distance predicate for a given set, whichever is easier to write down (although $0 \in I$ does not have a definable indicator function).

First note that the following formula is 0 if and only if $x = t$ and 1 otherwise

$$\chi(x) = \sup_y d(x, y) \wedge (4 \dot{-} d(x, y)) \wedge 1,$$

because t is the only point for which there is no x with $2 \leq d(x, t) \leq 3$. So t is definable and we can use it as a constant to define distance predicates for A and each $A \times \{n\}$,

$$d(x, A) = \inf_y d(x, y) + 2|d(y, t) - 5|, \qquad (\dagger)$$

$$d(x, A \times \{n\}) = \inf_y d(x, y) + 2|d(y, t) - (4 + 2^{-n-1})|. \qquad (\dagger)$$

These formulas work because of how we defined distances to t. $|d(y, t) - 5|$ and $|d(y, t) - (4 + 2^{-n-1})|$ roughly give the distances to A and $A \times \{n\}$ and then we plug those into formulas similar to the one used in the proof of Proposition 9.19 in [5] to get exact distance predicates.

For each $n < \omega$, there is a definable bijection from A to $A \times \{n\}$ given by

$$d(y, f_n(x)) = 2^{n+1}(d(x, y) \dot{-} 2), \qquad (\circ)$$

and so for any $n < \omega$, we can define P_n on A by

$$P_n(x, y) = 2^{n+1}(d(f_n(x), f_{n+1}(y)) \dot{-} 1),$$

So X is the required uniformly definable imaginary, which clearly preserves embeddings. The interpretation preserves (uniform) d-finiteness of types, lack of strong infinite dimensionality, local compactnes, and local finite dimensionality by the same argument as in the proof of Lemma 4.4, i.e. the inclusion maps are open isometries-up-to-scaling.

The advertised set of bad points is $\{\infty, *\} \cup \bigcup_{n<\omega}(*, n)$. Since this is a closed compact set of $\varnothing$-definable points, it is algebraic over $\varnothing$.

(ii) $T_{\mathcal{L}}$ is a theory in the language of metric spaces of diameter 5. By Lemma 3.2 and Propositions 4.2 and 5.6, we only need to construct $T_{\mathcal{L}}$ in the case where $\mathcal{L}$ has one sort and is 1-Lipschitz.

Recall that for a given formula φ there are axioms determined by φ that hold exactly when φ is the distance predicate of a definable set[5, Theorem 9.12]. Since these axioms are a fixed set of reduced sentences depending on φ, they're uniformly computable from φ. $T_{\mathcal{L}}$ contains axioms saying that $\inf_y d(x, y) + 5\chi(y)$ is the distance predicate of a definable set, and that the corresponding set is non-empty and has diameter 0. Let t denote the unique element of that definable set, for the sake of making the following axioms simpler to write down. Let $f : [0, 5] \to [0, 1]$ be a computable total continuous function whose zeroset is precisely $Z = \{4 + 2^{-n-1} : n < \omega\} \cup \{4, 5\}$. $T_{\mathcal{L}}$ has the following axioms

$$\sup_x f(d(x, t)),$$

$$\inf_x |d(x, t) - r|, \text{ for r} \in Z.$$

The first says that distances to t only take on values in Z. The second says that every distance in Z comes arbitrarily close to being attained (for isolated points in Z, i.e. everything except 4, this implies that the distance is exactly attained).

$T_{\mathcal{L}}$ also has axioms saying that the formulas (†) are distance predicates to definable sets, i.e. that A and $A \times \{n\}$ are definable sets. We also need axioms saying that the formula (∘) defines isometries-up-to-scaling between A and $A \times \{n\}$, i.e. if we let $DP_x[P(x)]$ be the sentence given in [5, Theorem 9.12] asserting that P(x) is a distance predicate for a definable set, then we have

$$\sup_{x \in A} DP_{y \in A \times \{n\}}[2^{n+1}(d(x,y) \dot{-} 2)],$$

$$\sup_{y \in A \times \{n\}} DP_{x \in A}[2^{n+1}(d(x,y) \dot{-} 2)],$$

where $DP_{x \in D}[P(x)]$ is DP relativized to the definable set D. Those two axioms state that the predicate $2^{n+1}(d(x,y) \dot{-} 2)$ gives a distance predicate on A for each element of $A \times \{n\}$ and vice versa. We also need

$$\sup_{x \in A} \inf_{y \in A \times \{n\}} 2^{n+1}(d(x,y) \dot{-} 2),$$

$$\sup_{y \in A \times \{n\}} \inf_{x \in A} 2^{n+1}(d(x,y) \dot{-} 2),$$

which say that the distance predicates always give non-empty definable sets. And we need

$$\sup_{x \in A} \sup_{y_0, y_1 \in A \times \{n\}} d(y_0, y_1) \dot{-} 4\left[2^{n+1}(d(x,y_0) \dot{-} 2) + 2^{n+1}(d(x,y_1) \dot{-} 2)\right],$$

$$\sup_{y \in A \times \{n\}} \sup_{x_0, x_1 \in A} d(x_0, x_1) \dot{-} 4\left[2^{n+1}(d(x_0,y) \dot{-} 2) + 2^{n+1}(d(x_1,y) \dot{-} 2)\right],$$

which say that the definable sets have diameter 0, i.e. that they are single points.[2]

And finally we need to actually assert that this function is an isometry-up-to-scaling:

[2] These two are just parameterized relativizations, i.e. $Q(x,y) = \sup_{z \in D(y)} P(x,z)$, which do work in general for families of uniformly definable sets but weren't treated in [5]. Likewise the first, third, and fifth axioms in this series are just instances of a schema

$$\sup_x DP_y[P(x,y)],$$

$$\sup_x \inf_x P(x,y),$$

$$\sup_x \sup_{y,z} d(y,z) \dot{-} 4[P(x,y) + P(x,z)],$$

asserting that $P(x,y)$ is a predicate defining a function.

$$\sup_{x_0,x_1\in A}\ \sup_{y_0,y_1\in A\times\{n\}} |d(x_0,x_1)-2^n d(y_0,y_1)| \dot{-} 2^{n+3}\left[(d(x_0,y_0)\dot{-}2)+(d(x_0,y_1)\dot{-}2)\right].$$

Furthermore we need axioms enforcing the definition of d given in part (i) of this proof, other than the line involving P_n (which isn't determined by $T_{\mathcal{L}}$) and lines involving ∞ (which are automatically enforced by continuity). The distances between A and $A\times\{n\}$ are already enforced by the previous axioms. We need the following

$$\sup_{x\in A}|d(x,t)-5|,$$

$$\sup_{x\in A\times\{n\}}\ \sup_{y\in A\times\{m\}} |d(x,y)-|2^{-n}-2^{-m}||, \text{ for } |n-m|>1,$$

$$\sup_{x\in A\times\{n\}}|d(x,t)-(4+2^{-n-1})|.$$

For the P_n line we just need to enforce the lower bound of 2^{-n-1}, the upper bound of 2^{-n}, and that P_n (which is definable from d since we can define the sets $A\times\{n\}$) obeys the correct modulus of uniform continuity (relative to the predicate $2P_0$, since the metric itself may have lost information due to clipping),

$$\sup_{x\in A\times\{n\}}\ \sup_{y\in A\times\{n+1\}} (2^{-n-1}\dot{-}d(x,y))\vee(d(x,y)\dot{-}2^{-n}),$$

$$\sup_{x_0,x_1,y_0,y_1\in A}|P_n(x_0,x_1)-P_n(y_0,y_1)|\dot{-}\Delta_{P_n}(2(P_0(x_0,y_0)\vee P_0(x_1,y_1))).$$

For those predicate symbols that were originally unary we need axioms enforcing that P_n only depends on one input,

$$\sup_{x,y_0,y_1\in A}|P_n(x,y_0)-P_n(x,y_1)|,\ P_n \text{ unary}.$$

The existence of ∞ and its definability are implied by these other axioms (since the $A\times\{n\}$ form a Cauchy sequence of definable sets in the Hausdorff metric whose diameters are limiting to 0 and a Hausdorff metric limit of definable sets is definable). Finally Ξ is just a sentence that evaluates to the diameter of the set A,

$$\Xi=\sup_{x,y\in A} d(x,y).$$

□

Assuming that the signature has finitely many sorts and a uniform arity bound (but maybe infinitely many predicate symbols) we can avoid the bad points entirely, but the construction is different. It is somewhat less delicate than the construction in Theorem 6.1, so we'll only sketch the important specifics.

Theorem 6.2. *If $\mathcal{L}$ is a countable metric signature with finitely many sorts and a uniform arity bound, then the result of Theorem 6.1 holds with no bad points, i.e. the bi-interpretation preserves local compactness and finite dimensionality everywhere.*

Proof. By applying Proposition 5.6, we may assume that $\mathcal{L}$ is 1-Lipschitz. Let $\{O_n\}_{n<k}$ be a finite list of all base sorts and let $\{N_n\}_{n<\ell}$ be a finite list of all finitary product sorts of the form $\prod_{O\in a(P)} O$ for some predicate symbol P. The sort X will be constructed from a graph with the following nodes:

- For each $n < k$, a main copy of the sort O_n.
- For each $n < \ell$, a copy of $N_n = \prod_{O\in a(P)} O$ along with copies of each O in $a(P)$ (with multiplicity).
- For each $n < \ell$, a copy of $I = \{0\} \cup \{2^{-s} : s < \omega\}$.

Connections between the nodes will correspond to specific relationships being encoded in the metric.

- For each main copy of O_n and each copy of O_n associated to some N_m there is an edge. Call the associated copy O'_n. The metric between $x \in O_n$ and $y \in O'_n$ will be given by $d(x, y) = 2 + d_{O_n}(x, y)$, in order to encode a definable bijection between O_n and O'_n.
- For each N_m and associated O'_n there is an edge. If O'_n is the ith factor of N_m, then the metric between $\overline{x} \in N_m$ and $y \in O'_n$ will be given by $d(\overline{x}, y) = 2 + d_{O_n}(x_i, y)$, in order to encode a definable projection from N_m to O'_n.
- For each N_m and its associated copy I_m of I there is an edge. Let $\{P_n\}_{n<\omega}$ be a list of the predicates symbols on N_m. If $\overline{x} \in N_m$ and $2^{-n} \in I$, then $d(\overline{x}, 2^{-n}) = 2 + 2^{-n}P_n(\overline{x})$ and $d(\overline{x}, 0) = 2$. (This is where it's important that the predicate symbols be 1-Lipschitz. If P_n is not 1-Lipschitz, this formula cannot define a metric).

Let all other distances be 4. Finally add a single new point t, with distances to everything else between 5 and 6 chosen to make each node of the graph have a t-definable indicator function. Then using the same kind of formula as in the proof of Theorem 6.1, t is $\varnothing$-definable, so each of the nodes in the graph is definable as well.

Note that for any function $\eta(x)$ taking on 0 on some copy of I and 1 everywhere else the formula

$$\left(\eta(x) + 8 \sup_y d(x, y) \wedge \left(\frac{1}{2} \dot{-} d(x, y)\right)\right) \wedge 1$$

is $\{0, 1\}$-valued and takes on the value 0 if and only if $x = 1 \in I$, so we can use $1 \in I$ as a constant and for each $n \leq \omega$, we can define a distance predicate for $2^{-n} \in I$ (with $2^{-\omega} = 0$) by

$$d(x, 2^{-n}) = \inf_y d(x, y) + 2|d(y, 1) - (1 - 2^{-n})|$$

So each point in each copy of I is $\varnothing$-definable.

Every point in $X^{\mathfrak{A}}$ is either an image of some Cartesian product of sorts in $\mathfrak{A}$ or contained in a compact clopen definable set (either a copy of I or t). Finitary products preserve local compactness and local finite dimensionality, so in this construction there are no 'bad points.' $\square$

6.1 Finite axiomatizability in continuous logic

The notion of finite axiomatizability is somewhat awkward in continuous logic. There are several possible definitions that suggest themselves, but none of them seem useful. This is the most literal transcription of the ordinary definition:

Definition 6.3 (Finite axiomatizability version 1)**.** A theory T is *finitely axiomatizable* if and only if it is axiomatized by a finite collection of sentences.

Depending on what we mean by 'sentence' every theory in a countable language is finitely axiomatizable in that continuous logic naturally has an infinitary conjunction of the form $\Sigma_{n<\omega} 2^{-n} \varphi_n$ and we can just let φ_n be an enumeration of a countable dense subset of the logical consequences of T.

A sensible attempt to avoid this would be a definition like this:

Definition 6.4 (Finite axiomatizability version 2)**.** A theory T is *finitely axiomatizable* if and only if it is axiomatized by a finite collection of restricted sentences.

But this is arbitrary and fails to have any obvious meaningful semantic consequences.

We can try a more directly semantic definition like this:

Definition 6.5 (Finite axiomatizability version 3)**.** A theory T is *finitely axiomatizable* if and only if the class of models of T and its complement are both elementary.

Which amounts to saying $[T] = \{T' \in S_0(\varnothing) : T' \vdash T, T' \text{ a complete theory}\}$ is a clopen subset of $S_0(\varnothing)$. The problem is that for any reasonable[3] metric signature $S_0(\varnothing)$ is connected, so the only finitely axiomatizable theories are the trivial theory and the inconsistent one. That said, 'finite axiomatizability version 3' relative to a theory can be non-trivial.

[3] If all function symbols in $\mathcal{L}$ have concave moduli of continuity, then $S_0(\varnothing)$ can be continuously retracted to a point by scaling all non-metric relations to 0 and then scaling the metric to 0. On the other hand, if $\Delta_f(x) = x^2$ and the metric has diameter ≤ 1, then the sentence

$$\inf_x d(x, f(x))$$

can only take on the values 0 or 1: Either there exists some x such that $d(x, f(x)) < 1$, in which case $x, f(x), f(f(x)), \ldots$ converges to a fixed point of f, or for every x, $d(x, f(x)) = 1$.

At this point we could argue that clopenness in type space is too strong of a condition in continuous logic. Definable sets do not correspond to clopen subsets of type space, but rather have a more subtle topometric characterization in terms of the d-metric: A closed set $D \subseteq S_n(T)$ is definable if and only if $D \subseteq \mathrm{int}\{p \in S_n(T) : d(p, D) < \varepsilon\}$ for every $\varepsilon > 0$, where $\mathrm{int}X$ is the topological interior of X. By analogy we could try a similar weakening of clopen as a basis for our definition of 'finite axiomatizability,' but the d-metric relies on T being a complete theory and for a complete theory $S_0(T)$ is trivial.

There are, however, contexts in which there is a meaningful non-trivial metric on $S_0(T)$ for an incomplete theory T, specifically if we're examining a notion of approximate isomorphism (such as the perturbations in [4] or Gromov-Hausdorff distance) we get a metric on completions of T:

$$\rho(T_0, T_1) = \inf\{\varepsilon : \mathfrak{A} \models T_0, \mathfrak{B} \models T_1, \mathfrak{A}, \mathfrak{B} \text{ `}\varepsilon\text{-isomorphic'}\},$$

whatever 'ε-isomorphic' might mean. And in this case we get a weaker notion of finite axiomatizability:

Definition 6.6 (Finite axiomatizability version 4)**.** A theory T is *finitely axiomatizable relative to* ρ if there is a sentence χ such that $T \vdash \chi$ and for all complete theories T', $T' \vdash \chi \leq \rho(T', [T])$, where $\rho(T', [T])$ is the point-set distance between T' and $[T]$.

Which is equivalent to the topometric condition $[T] \subseteq \mathrm{int}\{T' \in S_0(\varnothing) : \rho(T', [T]) < \varepsilon\}$ for every $\varepsilon > 0$. It should be noted that this is a proper generalization of version 3 in that we can take our notion of approximate isomorphism to be $\mathfrak{A}$ and $\mathfrak{B}$ are 0-isomorphic if they are isomorphic and 1-isomorphic if they are not.

This may be a reasonable definition in some context, although as discussed in [4] the metrics ρ are generally much more poorly behaved than the d-metric. In any case it's unclear what one can do with this definition and to apply it to this paper we would need to choose a notion of approximate isomorphism before we could even ask the question of whether or not the theory $T_{\mathcal{L}}$ is 'finitely axiomatizable.'

References

[1] Itaï Ben Yaacov and Alexander Usvyatsov. On d-finiteness in continuous structures. *Fundamenta Mathematicae*, 194(1):67–88, 0 2007.

[2] Itaï Ben Yaacov and Alexander Usvyatsov. Continuous first order logic and local stability. *Transactions of the American Mathematical Society*, 362(10):5213–5259, 2010.

[3] Itaï Ben Yaacov. Continuous first order logic for unbounded metric structures. *Journal of Mathematical Logic*, 8(2):197–223, 2008.

[4] Itaï Ben Yaacov. On perturbations of continuous structures. *Journal of Mathematical Logic*, 08(02):225–249, 2008.

[5] Itaï Ben Yaacov, Alexander Berenstein, C. Ward Henson, and Alexander Usvyatsov. *Model theory for metric structures*, volume 2 of *London Mathematical Society Lecture Note Series*, pages 315–427. Cambridge University Press, 2008.

[6] Wilfrid Hodges. *Model Theory*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1993.

[7] T. Kihara and A. Pauly. Point degree spectra of represented spaces. *ArXiv e-prints*, May 2014.

[8] Klaus Weihrauch. *Computable Analysis: An Introduction*. Springer-Verlag, Berlin, Heidelberg, 2000.