Journal of Logic & Analysis 17:1 (2025) 1–30
ISSN 1759-9008

Topometric Characterization of Type Spaces in Continuous Logic

JAMES E. HANSON

Abstract: We show that a topometric space X is topometrically isomorphic to a type space of some continuous first-order theory if and only if X is compact and has a metric satisfying that $\{p : d(p, U) < \varepsilon\}$ is open for every open U and $\varepsilon > 0$. Furthermore, we show that this can always be accomplished with a stable theory.

2020 Mathematics Subject Classification 03C66 (primary)

Keywords: continuous logic, topometric spaces

Introduction

Continuous first-order logic, introduced in Ben Yaacov, Berenstein, Henson and Usvyatsov [5], is a generalization of discrete (ie, classical two-valued) first-order logic suited for studying structures with natural metrics, such as C*–algebras (with various metrics such as the norm metric or the 2–distance induced by a normalized tracial state), valued fields, $\mathbb{R}$–trees, and probability algebras. Such structures are referred to as metric structures. In continuous logic, formulas take on arbitrary real values, arbitrary continuous functions are used as connectives, and supremum and infimum take on the role of quantifiers.

In continuous logic, there is important extra structure on type spaces, namely the induced metric given by:

$$\partial(p, q) \coloneqq \inf\{d(a, b) : a \models p, b \models q\}$$

This metric induces a topology on type space that is generally strictly finer than the compact logic topology. The induced metric enjoys certain strong compatibility properties with the logic topology, which Ben Yaacov isolated to develop a general theory of topometric spaces in [2].

Definition 0.1 A *topometric space* (X, τ, ∂) is a set X together with a topology τ and a metric ∂ such that the metric refines the topology and is lower semi-continuous (ie, has $\{(x, y) \in X^2 : \partial(x, y) \leq \varepsilon\}$ closed for every $\varepsilon > 0$).

Published: February 2025 https://doi.org/10.4115/jla.2025.17.1

A notably subtle aspect of the generalization to continuous logic is the treatment of definable sets. A subset D of a metric structure M is *definable* if $d(x,D) \coloneqq \inf_{y \in D} d(x,y)$ is a definable predicate.

In discrete logic, definable sets have a purely topological characterization in terms of clopen subsets of type space. Likewise, there is a purely topometric characterization of definable sets in continuous logic: A set[1] $D \subseteq S_n(T)$ corresponds to a definable set if and only if it is closed and has $D^{<\varepsilon} \coloneqq \{x \in S_n(T) : \partial(x,D) < \varepsilon\}$ open for every $\varepsilon > 0$. This means that the family of definable sets in a given type space can be read off from its topometric structure alone.

In [9], the author showed that there are strong restrictions on the kinds of type spaces that can occur in ω–stable (or more generally totally transcendental) theories. However, many counterexamples regarding definable sets occur in superstable theories (eg, Example 4.2 in [9], which is a weakly minimal theory T such that $S_1(T)$ has cardinality $2^{\aleph_0}$ but no non-trivial definable sets). These facts motivate a desire to solve the following questions: Which topometric spaces are type spaces of some continuous theory? What restrictions, if any, do various model-theoretic tameness conditions impose on the topometry of type spaces? Despite the existence of superstable theories that are poorly behaved with regards to definable sets, might there be, in principle, some subtler regularity imposed on them by stability or superstability?

In discrete logic, this question is fairly easy to answer:

Fact 0.2 *For any totally disconnected compact Hausdorff space X, there is a weakly minimal theory T such that $S_1(T) \cong X$. Furthermore, if X is scattered (ie, has ordinal Cantor–Bendixson rank), then T can be taken to be totally transcendental.*

Proof sketch Let $\mathcal{L}$ be a language consisting of a unary predicate P_K for each clopen set $K \subseteq X$. Let M be the $\mathcal{L}$–structure whose universe is X, where $P_K^M(x)$ holds if and only if $x \in K$. Let $\mathrm{Th}_{\mathbb{F}}(X)$ be the $\mathcal{L}$–theory of M. It is relatively straightforward to show that models $N \models \mathrm{Th}_{\mathbb{F}}(X)$ are characterized as sets N together with functions $f\colon N \to X$ with dense image satisfying that for any isolated $x \in X$, the fiber $f^{-1}(x)$ is a singleton. Furthermore, for any $x \in N$, $\mathrm{tp}(x)$ is uniquely determined by $f(x)$. This is enough to establish that $\mathrm{Th}_{\mathbb{F}}(X)$ has quantifier elimination and is weakly minimal. Furthermore, if X is scattered, it follows from quantifier elimination that $\mathrm{Th}_{\mathbb{F}}(X)$ is totally transcendental. $\square$

[1] By an abuse of terminology, we also refer to such sets of types as *definable*.

The metric ∂ on a given type space $S_n(T)$ has an additional regularity property not shared by compact topometric spaces in general, identified in Ben Yaacov [1] and referred to provisionally there as 'openness,' but later renamed to *adequacy* in Ben Yaacov and Melleray [6].

Definition 0.3 A topometric space (X, τ, ∂) has an *adequate* metric if for any open set U and any $\varepsilon > 0$, $U^{<\varepsilon}$ is an open set.

In this paper we will show that compactness and adequacy precisely characterize the topometric spaces that arise as type spaces in continuous logic. Furthermore, we will show that any compact topometric space with an adequate metric is topometrically isomorphic to a type space of a stable theory.

Although our result does resolve the general question completely, it still leaves open the characterization of the possible topometric structure of type spaces in totally transcendental and superstable theories.

Notation 0.4 In any topological or topometric space X, we will write $\operatorname{cl} A$ for the topological closure of the set A.

Note that for the purposes of this paper, closures in topometric spaces are always taken with regards to the topology rather than the metric.

The following lemma is an example of a slightly subtle consequence of adequacy (and compactness). It will allow us to sharpen our main result (in that we do not need to assume that ∂ is bounded) but does not fit neatly into the main body of the argument.

Lemma 0.5 *Any compact topometric space (X, τ, ∂) with adequate metric has finite diameter.*[2]

Proof If some non-empty open subset $U \subseteq X$ has finite diameter, then by compactness and adequacy there is some finite $\varepsilon > 0$ such that $X = U^{<\varepsilon}$, whereby X has finite diameter.

So assume for the sake of contradiction that every non-empty open subset of X has infinite diameter. Fix $x_1, y_1 \in X$ with $\partial(x_1, y_1) > 1$. By lower semi-continuity of ∂, there are open neighborhoods $U_1 \ni x_1$ and $V_1 \ni y_1$ such that $U_1 \times V_1$ is disjoint from

[2] Recall that for any set A in a metric space, the *diameter of* A, written $\operatorname{diam} A$, is $\sup\{d(a, a') : a, a' \in A\}$.

the closed set $\{(z, w) \in X^2 : \partial(z, w) \leq 1\}$. In particular, this means that $U_1^{<1}$ and V_1 are disjoint.

At stage i, given non-empty open sets U_i and V_i, since U_i and V_i both have infinite diameter, we can find $x_{i+1} \in U_i$ and $y_{i+1} \in V_i$ with $\partial(x_{i+1}, y_{i+1}) > i + 1$. Again by lower semi-continuity of ∂, we can find open neighborhoods $U_{i+1} \ni x_{i+1}$ and $V_{i+1} \ni y_{i+1}$ such that $\operatorname{cl} U_{i+1} \subseteq U_i$ and $\operatorname{cl} V_{i+1} \subseteq V_i$ and $U_{i+1}^{<i} \cap V_{i+1} = \varnothing$.

Let x_ω be an element of $\bigcap_{i<\omega} \operatorname{cl} U_i$ and y_ω of $\bigcap_{i<\omega} \operatorname{cl} V_i$, which are both non-empty by compactness. By construction we have that $\partial(x_\omega, y_\omega) > i$ for every $i < \omega$, but this contradicts that ∂ is a metric (rather than an extended metric). □

Finally, we should draw the reader's attention to the work of Carlisle and Henson on the model theory of $\mathbb{R}$–trees [7]. Although we do not use too many results from [7] directly, a fair number of the technical ideas used here are indebted to that paper. Moreover, we use many of the same references for result regarding $\mathbb{R}$–trees, such as Roe [10] and Chiswell [8].

1 Rationale and outline of the proof

It will be worthwhile to discuss for a moment why some more straightforward arguments seem to be insufficient to establish our main result. In other words, why is the proof in this paper so much more complicated than the proof of Fact 0.2? If we wanted to adapt the proof of Fact 0.2, the first thing to note is that we need to deal with continuous functions on X rather than clopen subsets of X, but this is an expected consideration in generalizing results from discrete logic to continuous logic. Moreover, by Ben Yaacov [3, Theorem 1.6], it is sufficient to only consider 1–Lipschitz functions from X to $\mathbb{R}$. If we were only concerned with the topology of the type space $S_1(T)$, this would be enough:

Fact 1.1 *For any compact Hausdorff space X, there is a weakly minimal continuous theory T such that $S_1(T) \cong X$. Furthermore, if X is scattered, then T can be taken to be totally transcendental.*

Proof sketch Repeat the argument in the proof of Fact 0.2 with a real-valued predicate P_f for each 1–Lipschitz function $f\colon X \to \mathbb{R}$. Give X a $\{0, 1\}$–valued metric with the obvious interpretations of the predicates P_f.[3] $\square$

Since ultimately our concern is with not just the topology of $S_1(T)$ but its full topometric structure, something more must be done. Given a compact topometric space (X, τ, ∂), we can think of this as a structure M_X with metric ∂ and unary predicates P_f for each 1–Lipschitz continuous function $f\colon X \to \mathbb{R}$. The issue is that even if ∂ is adequate, it can be the case that $S_1(\mathrm{Th}(M_X))$ is not even homeomorphic to X.

Example 1.2 *There is a compact topometric space (X, τ, ∂) with adequate metric such that $S_1(\mathrm{Th}(M_X))$ is not homeomorphic to X.*

Proof Let (X, τ) be the one-point compactification of $[0, 1] \times \mathbb{N}$. Let $*$ be the point added in the compactification. Define a metric ∂ on X by:

- $\partial((x, n), (y, n)) = |x - y|$
- $\partial((x, n), (y, m)) = 1$ for $n \neq m$
- $\partial((x, n), *) = 1$

It is relatively straightforward to verify that (X, τ, ∂) is a topometric space with an adequate metric.

Let $\varphi(x) = \sup_y \min\{d(x, y), 1 - d(x, y)\}$. Note that in the structure M_X, the point $*$ satisfies $\varphi(*) = 0$, but all other points (x, n) in M_X satisfy $\varphi((x, n)) = \frac{1}{2}$. Therefore $\mathrm{tp}(*)$ is an isolated point in $S_1(\mathrm{Th}(M_X))$. Since X has no isolated points, X and $S_1(\mathrm{Th}(M_X))$ are not homeomorphic. $\square$

This implies that we will need something more complicated. One possible approach would be a Fraïssé limit. Given (X, τ, ∂), call (M, d, g) an X–*structure* if (M, d) is a complete metric space (of diameter at most $\operatorname{diam} X$) and $g\colon M \to X$ is a 1–Lipschitz function. One would hope that the class of X–structures (understood as metric structures in an appropriate language containing the P_f's from before) is a metric Fraïssé class (in the sense of Ben Yaacov [4]) and that the resulting Fraïssé limit has the desired Stone space. While this should work, the issue with this approach is that the resulting

[3] One thing to note about the constructions in Facts 0.2 and 1.1 is that the behavior of the resulting structure is slightly non-uniform when X has isolated points. Specifically, the resulting theory will say for each isolated point of X that there is a unique element satisfying the associated type. We could have changed this by using a slightly more complicated definition of M_X. When we define $\mathbb{F}(X)$ in Definition 2.2, something similar occurs, but here the analogous issue is actually a problem and we need to adjust the construction accordingly. We discuss this more specifically after Definition 2.2.

theory would be quite model-theoretically wild. In particular, copies of the (scaled) Urysohn sphere would embed into the realizations of any given 1–type, implying that the resulting theory has TP_2 and SOP_n for every n.

In order to keep the theory model-theoretically tame, we would like to avoid as much higher arity structure as possible. One clean way to do this is with a tree-like structure, such as the $\mathbb{R}$–trees studied in [7], which are always stable. Since we would like to avoid choosing a root, it makes sense to consider '$\mathbb{R}$–forests' instead of $\mathbb{R}$–trees. Intuitively what we would like to be able to do is build some kind of 'generically X–colored $\mathbb{R}$–forests,' in which paths in the tree correspond to 1–Lipschitz paths in X. This too has a problem, which is that X may not be path-connected. Furthermore, even if X is path-connected, it may not admit 1–Lipschitz paths (eg, $[0,1]$ with the metric $d(x,y) = \sqrt{|x-y|}$).

This suggests that we need to allow 'jumps.' The theory we construct, $\mathrm{Th}_{\mathbb{F}}(X)$, is the theory of a particular structure $\mathbb{F}(X)$, which is meant to be generic among structures that admit embeddings into $\mathbb{R}$–forests with a 1–Lipschitz map into X. $\mathrm{Th}_{\mathbb{F}}(X)$ should be the model companion of this incomplete theory, but we have not verified this. We've decided to approach this construction in terms of building a specific model for two reasons: Firstly, we need to analyze the behavior of 'paths' in arbitrary models of our theory anyway (as we do in Section 4), so we wouldn't save much effort by starting with more combinatorially defined structures (such as discrete forests (R,E) with weighted edges and functions $f\colon R \to X$ that are 1–Lipschitz with regards to weighted path metric). Secondly, this way we have that all n–types with pairwise finite distances are realized in our explicitly constructed model $\mathbb{F}(X)$ (Corollary 5.10). This makes the analysis of arbitrary models of the theory $\mathrm{Th}_{\mathbb{F}}(X)$ easier, since we can verify desired behavior in $\mathbb{F}(X)$ directly instead of checking that various properties are formalizable in continuous first-order logic.

With that rationale in mind, the rest of the paper is structured as follows:

- In Section 2, we define the language and intended model $\mathbb{F}(X)$ of our theory for a given topometric space (X,τ,∂).
- In Section 3, we show that the finite-distance components of $\mathbb{F}(X)$ embed into $\mathbb{R}$–trees in a useful way.
- In Section 4, we analyze the collection $\mathcal{I}(X)$ of 'paths' in X (which we refer to as *itineraries* to avoid confusion with either the standard topological sense or the combinatorial sense of the word). In particular, we define a natural uniform structure (and therefore topology) on it which we then show is compatible with ultraproducts.

- In Section 5, we study other models of the theory $\mathrm{Th}_{\mathbb{F}}(X)$ of $\mathbb{F}(X)$ in order to fully characterize its types. We show that adequacy of the metric ensures that $S_1(\mathrm{Th}_{\mathbb{F}}(X))$ is isomorphic to X.
- In Section 6, we use the characterization of types in Section 5 to show that $\mathrm{Th}_{\mathbb{F}}(X)$ is strictly stable.
- Finally, in Section 7, we combine results into our main result, Theorem 7.1.

2 The forest-like structure associated to a topometric space

We will now define a language $\mathcal{L}_X$ and an $\mathcal{L}_X$–structure $\mathbb{F}(X)$ associated to a given compact topometric space (X, τ, ∂). Ultimately our goal is to show that $S_1(\mathrm{Th}(\mathbb{F}(X)))$ is topometrically isomorphic to X provided that ∂ is adequate.

Definition 2.1 Given a compact topometric space (X, τ, ∂) with finite positive diameter, we define an associated language $\mathcal{L}_X$ consisting of an $\mathbb{R}$–valued unary predicate P_f for each 1–Lipschitz function $f\colon X \to \mathbb{R}$, and a $\left(2 + \frac{r}{\operatorname{diam} X}\right)$–Lipschitz $[0, r]$–valued binary predicate d_r for each real $r > 0$. For each P_f, we let the associated modulus of uniform continuity be $\alpha_{P_f}(x) = x$. The 'official' metric of $\mathcal{L}_X$ is $d_{\operatorname{diam} X}$.

Definition 2.2 We write $\mathbb{F}(X)$ for the $\mathcal{L}_X$–structure whose universe consists of triples $K = (\operatorname{dom} K, \chi^K, \beta^K)$ where

- $\operatorname{dom} K$ is a compact subset of $\mathbb{R}_{\geq 0}$ containing 0,
- $\chi^K : \operatorname{dom} K \to X$ is a 1–Lipschitz function, and
- $\beta^K : (\operatorname{dom} K \setminus \{\sup \operatorname{dom} K\}) \to \omega$ is an arbitrary function.

For any $\mathcal{L}_X$–predicate P_f, we set:

$$P_f^{\mathbb{F}(X)}(K) = f(\chi^K(\sup \operatorname{dom} K))$$

We will define the metric(s) on $\mathbb{F}(X)$ in a moment with the help of the following ancillary notions:

Definition 2.3 We will write $\|K\|$ for $\sup \operatorname{dom} K$, which we will call the *length* of K.

For K and L in $\mathbb{F}(X)$, we say that L *extends* K, written $L \sqsupseteq K$, if $\operatorname{dom} L \supseteq \operatorname{dom} K$, $\chi^L \restriction \operatorname{dom} K = \chi^K$, and $\beta^L \restriction (\operatorname{dom} K \setminus \{\|K\|\}) = \beta^K$. For any $K \in \mathbb{F}(X)$ and $r \in [0, \|K\|]$, we write $K \restriction [0, r]$ for the unique maximal element L of $\mathbb{F}(X)$ such that $L \sqsubseteq K$ and $\|L\| \leq r$. We call elements of this form *initial segments* of K. As is typical, we will write $K \sqsupseteq L$ to mean $L \sqsubseteq K$.

For any $p \in X$, we write $\mathbb{F}(X,p)$ for the set $\{K \in \mathbb{F}(X) : \chi^K(0) = p\}$. We say that two elements K and K' of $\mathbb{F}(X)$ are in the same *finite-distance component* of $\mathbb{F}(X)$ (or that they have *finite distance*) if $\chi^K(0) = \chi^{K'}(0)$ (ie, if there is a $p \in X$ such that $K, K' \in \mathbb{F}(X,p)$). We sometimes refer to p as the *root*[4] of the finite-distance component $\mathbb{F}(X,p)$.

Finally for any K and K', either K and K' are not in the same finite-distance component or there is a unique largest r such that $r \in \operatorname{dom} K \cap \operatorname{dom} K'$ and $K\restriction[0,r] = K'\restriction[0,r]$. The element $K\restriction r = K'\restriction r$ is the *longest common initial segment of* K *and* K', written $K \sqcap K'$ if it exists.

Definition 2.4 We define an extended metric $d^{\mathbb{F}(X)}$ on $\mathbb{F}(X)$ by setting $d^{\mathbb{F}(X)}(K,K')$ equal to ∞ if K and K' are not in the same finite-distance component and $\|K\| + \|K'\| - 2\|K \sqcap K'\|$ otherwise.[5] We will write $d^{\mathbb{F}(X)}$ as d if no confusion can arise. For each $s > 0$, we set $d_s^{\mathbb{F}(X)} = \min\{d^{\mathbb{F}(X)}, s\}$.

Note that if $K \sqsubseteq K'$, then $d(K,K') = \|K'\| - \|K\|$.

The function β^K is included in Definition 2.2 to deal with the subtlety discussed in Footnote 3. In particular, we need it in Lemma 5.6. Without it the following issue could occur: Suppose we defined K as $(\operatorname{dom} K, \chi^K)$ and furthermore suppose that $p \in X$ is topologically isolated and has distance at least 1 from all other points of X. The metrically connected component of $(\{0\}, x \mapsto p)$ (where $x \mapsto p$ is the constant function with value p) would be elements of the form $([0,r], x \mapsto p)$ and would be isometric to $\mathbb{R}_{\geq 0}$. Furthermore, elements of this connected component would have distance at least 1 from all other elements of $\mathbb{F}(X)$. This means that in elementary extensions of $\mathbb{F}(X)$, the metrically connected component of $(\{0\}, x \mapsto p)$ will always be isometric to $\mathbb{R}_{\geq 0}$. Therefore $(\{0\}, x \mapsto p)$ and $([0,1], x \mapsto p)$ would not have the same type, but our goal is for the type of K to depend only on $\chi^K(\|K\|)$ (ie, the 'X–color' of the largest element of $\operatorname{dom} K$).

To fix this issue, the inclusion of β^K forces the structure to always have infinite branching at every point. For instance, if $\chi^K(0) = \chi^L(0) = p \in X$, $\|K\|$ and $\|L\|$ are both positive, and $\beta^K(0) \neq \beta^L(0)$, then $\|K \sqcap L\| = 0$ (ie, K and L represent 'paths' that have branched immediately). The use of ω is arbitrary. Any cardinal greater than 1 should produce an elementarily equivalent structure.

[4]This terminology is motivated by the fact that each $\mathbb{F}(X,p)$ is like an $\mathbb{R}$–tree in the same manner that $\mathbb{F}(X)$ as a whole is like an $\mathbb{R}$–forest. Note though that $\mathbb{F}(X,p)$ merely embeds into an $\mathbb{R}$–tree, and $\mathbb{F}(X)$ merely embeds into an $\mathbb{R}$–forest.

[5]Note that d is not part of the language $\mathcal{L}_X$.

Note that the domain of β^K is taken to exclude $\|K\|$ in order to avoid producing distinct K with distance 0.

Proposition 2.5 *In any $\mathbb{F}(X)$, d is a well-defined extended metric.*

Proof For any three K, K', and K'', if any two of them have infinite distance, then the triangle inequality is clearly satisfied, so assume that $K(0) = K'(0) = K''(0)$. $K \sqcap K'$ and $K' \sqcap K''$ must be comparable (since they're both initial segments of K'), so without loss of generality, we may assume that $\|K \sqcap K'\| \leq \|K' \sqcap K''\|$. We then necessarily have that $\|K \sqcap K''\| \geq \|K \sqcap K'\|$.

We now have that

$$d(K, K'') = \|K\| + \|K''\| - 2\|K \sqcap K''\|$$

and

$$d(K, K') + d(K', K'') = \|K\| + 2\|K'\| + \|K''\| - 2\|K \sqcap K'\| - 2\|K' \sqcap K''\|$$

so since $\|K' \sqcap K''\| \leq \|K'\|$ and $\|K \sqcap K'\| \leq \|K \sqcap K''\|$, we have that $2\|K \sqcap K'\| + 2\|K' \sqcap K''\| \leq 2\|K'\| + 2\|K \sqcap K''\|$ and the triangle inequality holds.

Finally, $d(x, y)$ is clearly symmetric and satisfies $d(x, y) = 0$ if and only if $x = y$. □

Note the easy fact that the absolute value $|\|K\| - \|K'\||$ is at most $d(K, K')$.

Lemma 2.6 *For any set $\mathfrak{S}$ of elements of $\mathbb{F}(X)$ with pairwise finite distance, there is a unique longest common initial segment $\bigsqcap \mathfrak{S}$ of $\mathfrak{S}$.*

Proof Let

$$r = \sup\{s : (\forall K \in \mathfrak{S}) s \in \pi(K) \wedge (\forall K, K' \in \mathfrak{S}) K{\upharpoonright}[0, s] = K'{\upharpoonright}[0, s]\}.$$

By continuity of χ^K and $\chi^{K'}$, we have that $(K{\upharpoonright}[0, r])_X = (K'{\upharpoonright}[0, r])_X$ for any $K, K' \in \mathfrak{S}$. Therefore, $K{\upharpoonright}[0, r] = K'{\upharpoonright}[0, r]$ for any $K, K' \in \mathfrak{S}$, and this is the required longest common initial segment. □

Lemma 2.7 *For any set $\mathfrak{S}$ of elements of $\mathbb{F}(X)$ with pairwise finite distance,*

$$\left\|\bigsqcap \mathfrak{S}\right\| \geq \sup_{K \in \mathfrak{S}} \|K\| - \operatorname{diam} \mathfrak{S}.$$

Proof Fix $K \in \mathfrak{S}$. Since all $K' \in \mathfrak{S}$ have $d(K,K') \leq \operatorname{diam}\mathfrak{S}$, we have:

$$\|K\| + \|K'\| - 2\|K \sqcap K'\| \leq \operatorname{diam}\mathfrak{S}$$
$$\frac{1}{2}(\|K\| + \|K'\| - \operatorname{diam}\mathfrak{S}) \leq \|K \sqcap K'\|$$

This implies that all elements of $\mathfrak{S}$ share with K a common initial segment of length at least

$$\frac{1}{2}(\|K\| + \inf_{K' \in \mathfrak{S}} \|K'\| - \operatorname{diam}\mathfrak{S})$$

which means that $\bigsqcap \mathfrak{S}$ is at least this long. Taking the supremum over $K \in \mathfrak{S}$ gives

$$\left\|\bigsqcap \mathfrak{S}\right\| \geq \frac{1}{2}\left(\sup_{K \in \mathfrak{S}} \|K\| + \inf_{K \in \mathfrak{S}} \|K\| - \operatorname{diam}\mathfrak{S}\right).$$

Clearly $\inf_{K \in \mathfrak{S}} \|K\| \geq \sup_{K \in \mathfrak{S}} \|K\| - \operatorname{diam}\mathfrak{S}$, so we have:

$$\left\|\bigsqcap \mathfrak{S}\right\| \geq \frac{1}{2}\left(2 \sup_{K \in \mathfrak{S}} - 2 \operatorname{diam}\mathfrak{S}\right) = \sup_{K \in \mathfrak{S}} \|K\| - \operatorname{diam}\mathfrak{S}$$ □

Corollary 2.8 *If $\mathfrak{S} \subseteq \mathbb{F}(X)$ has diameter at most $r < \infty$, then for any $K \in \mathfrak{S}$, $d(K, \bigsqcap \mathfrak{S}) \leq \operatorname{diam}\mathfrak{S}$.*

Proof For any $K \in \mathfrak{S}$, we have that

$$d\left(K, \bigsqcap \mathfrak{S}\right) = \|K\| - \left\|\bigsqcap \mathfrak{S}\right\| \leq \|K\| + \operatorname{diam}\mathfrak{S} - \sup_{K \in \mathfrak{S}} \|K\| \leq \operatorname{diam}\mathfrak{S}$$

as required. □

Proposition 2.9 *The metric d on $\mathbb{F}(X)$ is complete.*

Proof Let $\{K_i\}_{i<\omega}$ be a Cauchy sequence in $\mathbb{F}(X)$. By passing to a final segment, we may assume that the elements of this sequence have pairwise finite distance. Let $L_j = \bigsqcap\{K_i : i \geq j\}$. It is clear that L_j is an increasing sequence in the sense that $L_{j+1} \sqsupseteq L_j$ for every $j < \omega$. Furthermore, by Corollary 2.8, we have that $d(K_i, L_i) \to 0$ as $i \to \infty$.

Define $A = (\operatorname{dom} A, \chi^A, \beta^A)$ as follows: Let $\operatorname{dom} A = \operatorname{cl}\left(\bigcup_{i<\omega} \operatorname{dom} L_i\right)$, let χ^A be the unique 1–Lipschitz extension of $\bigcup_{i<\omega} \chi^{L_i}$, and let $\beta^A = \bigcup_{i<\omega} \beta^{L_i}$. By construction we have that A is the limit of the sequence $\{L_i\}_{i<\omega}$ and therefore of the sequence $\{K_i\}_{i<\omega}$ as well. □

Proposition 2.10 *For any 1–Lipschitz function $f\colon X \to \mathbb{R}$, the interpretation $P_f^{\mathbb{F}(X)}$ is 1–Lipschitz (with regards to the 'official' metric $d_{\operatorname{diam} X}$).*

Proof Fix $K, K' \in \mathbb{F}(X)$. If $d(K,K') \geq \operatorname{diam} X$, then there is nothing to prove, so assume that $d(K,K') < \operatorname{diam} X$. Since the induced functions χ^K and $\chi^{K'}$ are 1–Lipschitz, we have that:

$$\begin{aligned}\partial(\chi^K(\|K\|), \chi^{K'}(\|K'\|)) &\leq \partial(\chi^K(\|K\|), \chi^K(\|K \sqcap K'\|)) + \partial(\chi^{K'}(\|K \sqcap K'\|), \chi^{K'}(\|K'\|)) \\ &\leq d(K, K \sqcap K') + d(K \sqcap K', K') \\ &= d(K, K')\end{aligned}$$

Therefore:

$$\begin{aligned}\|P_f(K) - P_f(K')\| &\leq \partial(\chi^K(\|K\|), \chi^{K'}(\|K\|)) \\ &\leq \min\{d(K,K'), \operatorname{diam} X\} \\ &= d_{\operatorname{diam} X}(K, K')\end{aligned}$$ □

Corollary 2.11 $\mathbb{F}(X)$ *is an* $\mathcal{L}_X$*–structure.*

Proof Given Proposition 2.10, the only thing to verify is that the predicate symbols d_r obey the correct moduli of uniform continuity relative to the 'official' metric $d_{\operatorname{diam} X}$. For any K, K', L, and L', we have that:

$$\begin{aligned}\|d_r(K,L) - d_r(K',L')\| &\leq \min\{\|d(K,L) - d(K',L')\|, r\} \\ &\leq \min\{2\max\{d(K,K'), d(L,L')\}, r\} \\ &\leq \left(2 + \frac{r}{\operatorname{diam} X}\right) \min\{\max\{d(K,K'), d(L,L')\}, \operatorname{diam} X\} \\ &= \left(2 + \frac{r}{\operatorname{diam} X}\right) \max\{d_{\operatorname{diam} X}(K,K'), d_{\operatorname{diam} X}(L,L')\}\end{aligned}$$ □

Definition 2.12 For any X, let $\mathrm{Th}_{\mathbb{F}}(X)$ be $\mathrm{Th}(\mathbb{F}(X))$.

In any model M of $\mathrm{Th}_{\mathbb{F}}(X)$, we define an extended metric d by setting $d(x,y) = \sup_r d_r(x,y)$. The theory of $\mathbb{F}(X)$ ensures that this is an extended metric satisfying $d_r(x,y) = \min\{d(x,y), r\}$ for every $r > 0$.

3 $\mathbb{R}$–tree embeddings of components of $\mathbb{F}(X)$

Here we will show that the finite-distance components of $\mathbb{F}(X)$ embed isometrically into $\mathbb{R}$–trees in a particularly compatible way. This will be useful later in Section 5 when verifying certain properties of $\mathrm{Th}_{\mathbb{F}}(X)$. We will also take the opportunity to review some relevant properties of $\mathbb{R}$–trees.

Definition 3.1 A metric space (Z,d) is *geodesic* if for any $x,y\in Z$, there is an isometric path $f:[0,d(x,y)]\to Z$ with $f(0)=x$ and $f(d(x,y))=y$. In a geodesic space Z, we'll write $[x,y]$ for the image of some such path from x to y. (In our context, $[x,y]$ will always be unique.)

An $\mathbb{R}$–*tree* is a geodesic metric space that is uniquely arc-connected.

A metric space (Z,d) is 0–*hyperbolic* if it satisfies the 4–*point condition*: for all $x,y,z,w\in Z$,

$$d(x,y)+d(z,w)\le\max\{d(x,z)+d(y,w),d(y,z)+d(x,w)\}.$$

Note that it is immediate that if Z is an $\mathbb{R}$–tree and $\{W_i\}_{i\in I}$ is any family of sub-$\mathbb{R}$–trees of Z (ie, subspaces that are also $\mathbb{R}$–trees), then $\bigcap_{i\in I}W_i$ is also an $\mathbb{R}$–tree. In particular, this implies that any subset of W an $\mathbb{R}$–tree Z is contained in a unique minimal $\mathbb{R}$–tree. Following [7, Definition 2.12], we'll denote this sub-$\mathbb{R}$–tree E_W.

Note that $[x,y]=E_{\{x,y\}}$ and $E_W=\bigcup_{x,y\in W}[x,y]$. Furthermore, note that $\operatorname{cl}E_W$ is also a sub-$\mathbb{R}$–tree for any $W\subseteq Z$.

Definition 3.2 For any points x, y, and z, the *Gromov product of y and z at x* is

$$(y|z)_x=\frac{1}{2}(d(x,y)+d(x,z)-d(y,z)).$$

The following facts are mentioned in Carlisle and Henson [7] in Lemma 2.14, after Definition 2.8, and after Lemma 2.5.

Fact 3.3 *Fix a metric space (Z,d).*

1. *Z is 0–hyperbolic if and only if it isometrically embeds into an $\mathbb{R}$–tree [10, Proposition 6.12].*
2. *Any subset of an $\mathbb{R}$–tree is contained in a unique minimal sub-$\mathbb{R}$–tree.*
3. *Assume Z is 0–hyperbolic and $f_i:Z\to W_i$ are isometric embeddings into $\mathbb{R}$–trees for both $i<2$. If E_i is the unique minimal sub-$\mathbb{R}$–tree of W_i containing $f_i(Z)$ for both $i<2$, then there is a unique isometry $g:E_0\to E_1$ satisfying that $g\circ f_1=f_2$ [7, Lemma 2.14 (3)].*
4. *Z is an $\mathbb{R}$–tree if and only if it is geodesic and 0–hyperbolic [10, Proposition 6.12].*
5. *If Z is an $\mathbb{R}$–tree, then for any $x,y,z\in Z$, $d(x,[y,z])=(y|z)_x$ [10, Corollary 6.9].*
6. *If Z is an $\mathbb{R}$–tree, then for any closed sub-$\mathbb{R}$–tree $W\subseteq Z$ and any $z\in Z$, there is a unique $w\in W$ satisfying $d(z,w)=d(z,W)$ [8, Lemma 2.1.9].*

In particular, the last part of Fact 3.3 implies that if Z is an $\mathbb{R}$–tree, then for any $x, y, z \in Z$, there is a unique point $w \in [y, z]$ satisfying that $d(x, [y, z]) = d(x, w)$. We will follow [7, Note 2.15] and use the notation $Y(x, y, z)$ for this point w. Note that this function is symmetric: $Y(x, y, z) = Y(y, x, z) = Y(z, x, y)$.

Now we apply these facts to $\mathbb{F}(X)$.

Proposition 3.4 *For each $p \in X$, the finite-distance component $F(X, p)$ of $\mathbb{F}(X)$ is 0–hyperbolic and therefore isometrically embeds into an $\mathbb{R}$–tree.*

Proof Fix $p \in \mathbb{F}(X)$ and consider the finite-distance component $\mathbb{F}(X, p)$. Fix $K_0, K_1, K_2, K_3 \in \mathbb{F}(X, p)$. Let $d_{ijk\ell} = d(K_i, K_j) + d(K_k, K_\ell)$, and let $K_{ij} = K_i \sqcap K_j$. We need to verify that $d_{0123} \leq \max\{d_{0213}, d_{1203}\}$. Expanding definitions, we have that:

$$\begin{aligned} d_{ijk\ell} &= \|K_i\| + \|K_j\| - 2\|K_{ij}\| + \|K_k\| + \|K_\ell\| - 2\|K_{k\ell}\| \\ &= \|K_i\| + \|K_j\| + \|K_k\| + \|K_\ell\| - 2(\|K_{ij}\| + \|K_{k\ell}\|) \end{aligned}$$

Therefore, $d_{0123} \leq \max\{d_{0213}, d_{1203}\}$ is equivalent to:

$$\begin{aligned} -2(\|K_{01}\| + \|K_{23}\|) &\leq \max\{-2(\|K_{02}\| + \|K_{13}\|), -2(\|K_{12}\| + \|K_{03}\|)\}, \\ \|K_{01}\| + \|K_{23}\| &\geq \min\{\|K_{02}\| + \|K_{13}\|, \|K_{12}\| + \|K_{03}\|\} \end{aligned}$$

Let $K_{0123} = K_0 \sqcap K_1 \sqcap K_2 \sqcap K_3$. Up to symmetry, there are only 5 (non-exclusive) cases:

- $K_{01} \sqsupseteq K_{02} = K_{12} \sqsupseteq K_{03} = K_{13} = K_{23} = K_{0123}$, in which case $\|K_{01}\| \geq \|K_{02}\|$ and $\|K_{23}\| = \|K_{13}\|$.
- $K_{02} \sqsupseteq K_{01} = K_{21} \sqsupseteq K_{03} = K_{23} = K_{13} = K_{0123}$, in which case $\|K_{01}\| = \|K_{12}\|$ and $\|K_{23}\| = \|K_{03}\|$.
- $K_{02} \sqsupseteq K_{03} = K_{23} \sqsupseteq K_{01} = K_{12} = K_{13} = K_{0123}$, in which case $\|K_{01}\| = \|K_{12}\|$ and $\|K_{23}\| = \|K_{03}\|$.
- $K_{02} = K_{03} = K_{12} = K_{13} = K_{0123}$, in which case $\|K_{01}\| \geq \|K_{02}\|$ and $\|K_{23}\| \geq \|K_{13}\|$.
- $K_{01} = K_{03} = K_{12} = K_{23} = K_{0123}$, in which case $\|K_{01}\| = \|K_{03}\|$ and $\|K_{23}\| = \|K_{12}\|$.

So in every case the 4–point condition holds. Therefore $(\mathbb{F}(X, p), d)$ is 0–hyperbolic and isometrically embeds into an $\mathbb{R}$–tree. □

Definition 3.5 For any $p \in X$, let $\mathbb{R}\mathrm{T}(X, p)$ be the unique minimal $\mathbb{R}$–tree extending the finite-distance component $\mathbb{F}(X, p)$.

For the moment we need to be careful about distinguishing between certain objects in $\mathbb{R}\mathrm{T}(X,p)$ and similar objects in $\mathbb{F}(X,p)$. To that end, given $a,b,c \in \mathbb{R}\mathrm{T}(X,p)$, we'll write $Y^{\mathbb{R}\mathrm{T}(X,p)}(a,b,c)$ for $Y(a,b,c)$ evaluated in $\mathbb{R}\mathrm{T}(X,p)$. Likewise, we'll write $[a,b]^{\mathbb{R}\mathrm{T}(X,p)}$ for $E_{\{a,b\}}$ evaluated in $\mathbb{R}\mathrm{T}(X,p)$ (ie, the unique geodesic in $\mathbb{R}\mathrm{T}(X,p)$ from a to b). The following proposition will allow us to be more cavalier with this notation.

Proposition 3.6 *For any $p \in X$ and $K_0,K_1,K_2 \in \mathbb{F}(X,p)$, $Y^{\mathbb{R}\mathrm{T}(X,p)}(K_0,K_1,K_2) \in \mathbb{F}(X,p)$.*

Proof By the symmetry of Y and by relabeling if necessary, we may assume that $K_0 \sqcap K_1 \sqsupseteq K_0 \sqcap K_2 = K_1 \sqcap K_2$. It is immediate by the definition of d that $Y^{\mathbb{R}\mathrm{T}(X,p)}(K_0,K_1,K_2) = K_0 \sqcap K_1$, which is an element of $\mathbb{F}(X,p)$. □

Note that $Y(x,y,z)$ is not literally a definable function in the sense of continuous logic, since it is only well-defined inside finite-distance components, which are co-type-definable. It is, however, representable as a family of partial definable functions with domains of the form $\{(x,y,z) : d(x,y), d(x,z), d(y,z) < r\}$ for each $r > 0$ (see Corollary 3.11).

Definition 3.7 Given $p \in X$ and $K_0,K_1,K_2 \in \mathbb{F}(X,p)$, we will write $Y(K_0,K_1,K_2)$ for the element of $\mathbb{F}(X,p)$ guaranteed by Proposition 3.6. We will also write $[K_0,K_1]$ for $[K_0,K_1]^{\mathbb{R}\mathrm{T}(X,p)} \cap \mathbb{F}(X,p)$.

Note that $[K_0,K_1]$ will in general not be connected, although it will always be isometric to a closed subset of $[0,d(K_0,K_1)]$.

Proposition 3.8 *Fix $p \in X$. For any $K,K',L \in \mathbb{F}(X,p)$, $L \in [K,K']$ if and only if either $K \sqcap K' \sqsubseteq L \sqsubseteq K$ or $K \sqcap K' \sqsubseteq L \sqsubseteq K'$.*

Proof We know by Fact 3.3 that $L \in [K,K']^{\mathbb{R}\mathrm{T}(X,p)}$ if and only if $(K|K')_L = 0$, ie, if and only if $d(L,K)+d(L,K') = d(K,K')$. If $K \sqcap K' \sqsubseteq L \sqsubseteq K$ or $K \sqcap K' \sqsubseteq L \sqsubseteq K'$, then $d(L,K)+d(L,K') = d(K,K')$ by the definition of d.

Assume that $K \sqcap K' \sqsubseteq L \sqsubseteq K$ and $K \sqcap K' \sqsubseteq L \sqsubseteq K'$ both fail. In order for this to happen, it must be the case that either

- L properly extends K,
- L properly extends K',
- $K \sqcap L$ is strictly $\sqsubseteq$–between $K \sqcap K'$ and K,
- $K' \sqcap L$ is strictly $\sqsubseteq$–between $K \sqcap K'$ and K', or

- $K \sqcap K' \sqcap L$ is a strict initial segment of $K \sqcap K'$.

In each of these cases we clearly get that $d(L,K) + d(L,K') > d(K,K')$. □

Proposition 3.9 *For any $r > 0$ and K and K' in $\mathbb{F}(X)$ with $d(K,K') < r$, the formula*

$$\delta_{K,K',r}(x) := \min\left\{\frac{1}{2}(d_{3r}(x,K) + d_{3r}(x,K') - d_r(K,K')), r\right\}$$

defines the distance predicate (with regards to the metric d_r) of $[K,K']$.

Proof We know from Fact 3.3 and Proposition 3.6 that for any K, K', and L in the same finite-distance component of $\mathbb{F}(X)$, $d(L,[K,K']) = (K|K')_L = \frac{1}{2}(d(L,K) + d(L,K') - d(K,K'))$. Since $d(K,K') = d_r(K,K')$, we know that if $d(L,K) \leq 3r$ and $d(L,K') \leq 3r$, then $\delta_{K,K',r}(L) = \min\{(K|K')_L, r\} = \min\{d(L,[K,K']), r\} = d_r(L,[K,K'])$.

So assume that $d(L,K) > 3r$. By the reverse triangle inequality, it must be the case that $d(L,A) > 2r$ for every $A \in [K,K']$. Therefore $d(L,[K,K']) \geq 2r$. Furthermore, $d(L,K') \geq 2r$, so $\delta_{K,K',r}(L) = r = d_r(L,[K,K'])$, as required. If $d(L,K') > 3r$, the same argument establishes that $\delta_{K,K',r}(L) = d_r(L,[K,K'])$. □

Proposition 3.9 establishes that for any $M \models \mathrm{Th}_{\mathbb{F}}(X)$ and any $a,b \in M$ with $d(a,b) < r$, $\delta_{a,b,r}(x)$ is also the distance predicate of a definable set.

Definition 3.10 For any $M \models \mathrm{Th}_{\mathbb{F}}(X)$ and any $a,b \in M$ with $d(a,b) < r$, we write $[a,b]^M$ for the definable set with distance predicate $\delta_{a,b,r}(x)$. If no confusion can arise, we may write this as $[a,b]$.

Corollary 3.11 *For any $K_0, K_1, K_2, L \in \mathbb{F}(X)$ with $d(K_i,K_j) < r$ for all $i < j < 3$, we have that*

$$d_r(L, Y(K_0,K_1,K_2)) = \max_{i<j<3} \delta_{K_i,K_2,r}(L).$$

Moreover, for any model $M \models \mathrm{Th}_{\mathbb{F}}(X)$ and $a_0,a_1,a_2 \in M$ with $d(a_i,a_j) < r$ for all $i < j < 3$, the formula $\max_{i<j<3} \delta_{a_i,a_j,r}(x)$ also defines a singleton.

Proof The first statement is immediate from Proposition 3.9. The second statement follows from the fact that 'if $d(a_i,a_j) < r$ for all $i < j < 3$, then $\max_{i<j<3} \delta_{a_i,a_j,r}(x)$ is the d_r–distance predicate of a singleton' is expressible in continuous first-order logic. □

Note that in general the intersection of definable sets is not definable, so the fact that the formula in Corollary 3.11 works relies on special properties of the geometry of 0–hyperbolic spaces.

Definition 3.12 For any $M \models \mathrm{Th}_{\mathbb{F}}(X)$ and any $a, b, c \in M$ with pairwise finite distance, we write $Y^M(a, b, c)$ for the unique element defined by the formula in Corollary 3.11. If no confusion can arise, we may write this as $Y(a, b, c)$.

Definition 3.13 In any model M of $\mathrm{Th}_{\mathbb{F}}(X)$, a *finite tree* is a set which can be written as a union of a finite sequence $\{[a_i, b_i]\}_{i<n}$ with the property that for each $i < n$ with $i > 0$, $[a_i, b_i]$ is not disjoint from $\bigcup_{j<i}[a_j, b_j]$.

Since a finite tree is a finite union of definable sets, it is itself always a definable set. Note also that all elements of a finite tree are automatically in the same finite-distance component. It is immediate from Fact 3.3 that for any finite tree R in $\mathbb{F}(X)$, $R = E_R \cap \mathbb{F}(X)$. Furthermore, since the finite-distance components of models of $\mathrm{Th}_{\mathbb{F}}(X)$ are 0–hyperbolic, similar statements are true of finite trees in arbitrary models as well.

Lemma 3.14 *If R is a finite tree, then for any $a, b \in R$, $[a, b] \subseteq R$.*

Proof Since sets of the form $[a, b]$ (with $d(a, b)$ bounded by a fixed r) are uniformly definable, it is enough to verify this in $\mathbb{F}(X)$. In $\mathbb{F}(X)$, this follows immediately from the fact that $R = E_R \cap \mathbb{F}(X)$ and for any $K, L \in R$, $[K, L] = [K, L]^{\mathbb{R}\mathrm{T}(X,p)} \cap \mathbb{F}(X)$ (where $\mathbb{F}(X, p)$ is the finite-distance component of R). □

Definition 3.15 For any model $M \models \mathrm{Th}_{\mathbb{F}}(X)$ and any finite tuple $\bar{a} \in M$ with pairwise finite distance, the *convex closure of* $\bar{a}$, written $\mathrm{ccl}(\bar{a})$, is the intersection of all finite trees containing $\bar{a}$.

We will not need to prove this, but for finite tuples $\bar{a}$ with pairwise finite distance, $\mathrm{ccl}(\bar{a})$ is actually both the definable and algebraic closures of $\bar{a}$. Furthermore, in $\mathbb{F}(X)$, $\mathrm{ccl}(\bar{a}) = E_{\bar{a}} \cap \mathbb{F}(X)$ for any such tuple $\bar{a}$ (where $E_{\bar{a}}$ is the minimal sub-$\mathbb{R}$–tree of $\mathbb{R}\mathrm{T}(X, p)$ containing $\bar{a}$ and p is the root of the finite-distance component of $\bar{a}$).

Note that for any finite tree $R = \bigcup_{i<n}[a_i, b_i]$, we have that $R = \mathrm{ccl}(a_0 b_0 \ldots a_{n-1} b_{n-1})$.

Proposition 3.16 *For any model M of $\mathrm{Th}_{\mathbb{F}}(X)$, any finite tree R in M, and any $b \in M$ in the same finite-distance component of M, there is a unique element $c \in R$ with $d(b, c) = d(b, R)$.*

Proof Let $R = \mathrm{ccl}(\bar{a})$. Since R is a finite union of sets of the form $[a_i, a_j]$, there must be some i and j such that $d(b, R) = d(b, [a_i, a_j])$. We then have that $Y(a_i, a_j, b) \in [a_i, a_j]$ is the required element of R. □

4 The space of itineraries

In this section we develop some machinery needed to analyze sets of the form $[a,b]$ in arbitrary models of $\mathrm{Th}_{\mathbb{F}}(X)$. One of our goals is to establish that every such set already occurs up to isomorphism in $\mathbb{F}(X)$. The set we are about to define, $\mathcal{I}(X)$, is meant to be the set of isomorphism types of sets of the form $[a,b]$ in models of $\mathrm{Th}_{\mathbb{F}}(X)$ (where a is taken to be a designated element, so that $[a,b]$ and $[b,a]$ are not necessarily isomorphic in the sense that we are considering).

Definition 4.1 For any compact topometric space (X,τ,∂), let the *set of itineraries in* X, written $\mathcal{I}(X)$, be the collection of all 1–Lipschitz functions $f\colon D\to X$ with compact domain $D\subseteq\mathbb{R}_{\geq 0}$ containing 0. We write $\|f\|$ for $\sup\operatorname{dom} f$.

For any $r\geq 0$, let $\mathcal{I}_r(X)$ be the set $\{f\in\mathcal{I}(X) : \|f\|\leq r\}$.

Note that for any $K\in\mathbb{F}(X)$, we have that χ^K is an element of $\mathcal{I}(X)$. Furthermore, every element of $\mathcal{I}(X)$ occurs in this way. Also note that elements of $\mathcal{I}(X)$ are automatically *topologically* continuous as well, since the ordinary topology is coarser than the metric topology.

We can also regard elements of $\mathcal{I}(X)$ as being $\mathcal{L}_X(c)$–structures (where c is an added constant symbol) in a natural way:

Definition 4.2 For any $f\in\mathcal{I}(X)$, let $\mathbb{I}(f)$ be the $\mathcal{L}_X(c)$–structure whose universe is $\operatorname{dom} f$ with the interpretation

- $c^{\mathbb{I}(f)}=0$,
- $d_r^{\mathbb{I}(f)}(x,y)=\min\{|x-y|,r\}$ for all $r\in\mathbb{R}$, and
- $P_g^{\mathbb{I}(f)}(x)=g(f(x))$ for all 1–Lipschitz $g:X\to\mathbb{R}$.

For each $r\geq 0$, let $\mathrm{Th}(\mathcal{I}_r(X))$ be the $\mathcal{L}_X(c)$–theory of the class of all $\mathbb{I}(f)$ for $f\in\mathcal{I}_r(X)$ (ie, $\mathrm{Th}(\mathcal{I}_r(X))$ is the set of closed $\mathcal{L}_X(c)$–conditions that hold in all $\mathbb{I}(f)$).

It is fairly immediate that for any $f\in\mathcal{I}(X)$, $\mathbb{I}(f)$ is a compact $\mathcal{L}_X$–structure.

Note that for any $K\in\mathbb{F}(X)$, the substructure of $\mathbb{F}(X)$ consisting of elements of the form $K\restriction[0,r]$ is isomorphic to (the $\mathcal{L}_X$–reduct of) $\mathbb{I}(\chi^K)$ with $c^{\mathbb{I}(\chi^K)}=K\restriction[0,0]$ (ie, the root of the finite-distance component of K).

Since, for each r, the models of $\mathrm{Th}(\mathcal{I}_r(X))$ are uniformly compact, we expect that ultraproducts will induce a compact Hausdorff topology on the set of isomorphism types of models of $\mathrm{Th}(\mathcal{I}_r(X))$. Our goal right now is to characterize this topology directly and show that every model of $\mathrm{Th}(\mathcal{I}_r(X))$ is $\mathbb{I}(f)$ for some $f \in \mathcal{I}_r(X)$. Since the isomorphism type of $\mathbb{I}(f)$ is uniquely determined by f, this will allow us to identify $\mathcal{I}_r(X)$ with the space of completions of $\mathrm{Th}(\mathcal{I}_r(X))$. We will later use this to characterize the types in $\mathrm{Th}_{\mathbb{F}}(X)$. To this end, we will put a uniform structure on $\mathcal{I}(X)$ whose restrictions to each $\mathcal{I}_r(X)$ will ultimately correspond to the natural compact topology on the space of completions of $\mathrm{Th}(\mathcal{I}_r(X))$.

Our uniform structure on $\mathcal{I}(X)$ is generated by entourages of the form $E_{V,\varepsilon}$ for V, an entourage[6] in X^2, and $\varepsilon > 0$, where $(f,g) \in E_{V,\varepsilon}$ if and only if

- for every $r \in \operatorname{dom} f$, there is an $s \in \operatorname{dom} g$ such that $(f(r), g(s)) \in V$ and $|r - s| < \varepsilon$ and
- for every $s \in \operatorname{dom} g$, there is an $r \in \operatorname{dom} f$ such that $(f(r), g(s)) \in V$ and $|r - s| < \varepsilon$.

To see that this generates a uniform structure on $\mathcal{I}(X)$, note that

- $E_{V\cap W,\min\{\varepsilon,\delta\}} \subseteq E_{V,\varepsilon} \cap E_{W,\delta}$ and
- if $W^{\circ 2} \coloneqq \{(x,z) : (\exists y \in X)(x,y) \in W \wedge (y,z) \in W\} \subseteq V$, then $E^{\circ 2}_{W,\varepsilon/2} \subseteq E_{V,\varepsilon}$.

Recall that a uniform structure is *complete* if every Cauchy net converges, where a *Cauchy net* is a net $\{x_i\}_{i\in I}$ such that for every entourage $V \subseteq X^2$, there is an $i \in I$ such that $(x_j, x_k) \in V$ for all $j, k \geq i$. A uniform structure is *Hausdorff* if the induced topology is Hausdorff. For any $V \subseteq X^2$, we write $V(x)$ for the set $\{y \in X : (x,y) \in V\}$.

Proposition 4.3 *The uniform structure on $\mathcal{I}(X)$ is Hausdorff and complete. Furthermore, each $\mathcal{I}_r(X)$ is closed.*

Proof First to see that the uniform structure on $\mathcal{I}(X)$ is Hausdorff, let f and g be distinct elements of $\mathcal{I}(X)$. If $\operatorname{dom} f = \operatorname{dom} g$, then there must be an $\varepsilon > 0$ small enough that $(f,g) \notin E_{V,\varepsilon}$ for any entourage $V \subseteq X^2$. If $\operatorname{dom} f \neq \operatorname{dom} g$, then there must be some $r \in \operatorname{dom} f$ such that $f(r) \neq g(r)$. Find an entourage V small enough that $g(r) \notin \operatorname{cl} V(f(r))$, and then find $\varepsilon > 0$ small enough that $g(r) \notin (\operatorname{cl} V(f(r)))^{\leq\varepsilon}$. Now assume that $(f,g) \in E_{V,\varepsilon}$. By definition, this means that there is some $s \in \operatorname{dom} g$ with $|r-s| < \varepsilon$ such that $(f(r), g(s)) \in V$, ie, $g(s) \in V(f(r))$. Since g is 1–Lipschitz, we have that $\partial(g(r), g(s)) \leq |r - s| < \varepsilon$, which implies that $g(r) \in V(f(r))^{<\varepsilon} \subseteq (\operatorname{cl} V(f(r)))^{\leq\varepsilon}$, which is a contradiction. Therefore $(f,g) \notin E_{V,\varepsilon}$.

To show that the uniform structure is complete, let $\{f_i\}_{i\in I}$ be a Cauchy net on some directed set I.

[6]Recall that for compact spaces, there is a unique uniform structure compatible with the topology. In this case, $V \subseteq X^2$ is an entourage if and only if it is a neighborhood of the diagonal $\{(x,x) : x \in X\}$.

Let F be the set of points r in $\mathbb{R}_{\geq 0}$ with the property that for every $\varepsilon > 0$, there is an $i \in I$ such that for all $j \geq i$, r has distance at most ε from the domain of f_j. It is clear that $0 \in F$ and that F is closed. By looking at f_i for some sufficiently large $i \in I$, we can see that F must be bounded and therefore compact.

For each $r \in F$ and $i \in I$, let s_i^r be the smaller of the (one or two) nearest points in $\operatorname{dom} f_i$ to r. (Note that this is well defined since $\operatorname{dom} f_i$ is always non-empty.) Consider the net $\{f_i(s_i^r)\}_{i \in I}$ of points in X.

Claim. For each $r \in F$, the net $\{f_i(s_i^r)\}_{i \in I}$ is convergent.

Proof of claim. Fix an entourage $V \subseteq X^2$. Find an entourage $W \subseteq X^2$ and an $\varepsilon > 0$ small enough that the set

$$A \coloneqq \{(x, y) \in X^2 : \exists z (x, z) \in \operatorname{cl} W \wedge \partial(z, y) < 3\varepsilon\}$$

is contained in V. (This is always possible by compactness.) Now find $i \in I$ large enough that for any $j, k \geq i$, $(f_j, f_k) \in U_{W,\varepsilon}$ and the distance between r and the domain of f_j and f_k is at most ε. This implies that for any $j, k \geq i$, there is some $t \in \operatorname{dom} f_j$ such that $|s_j^r - t| < \varepsilon$ and $(f_j(s_j^r), f_k(t)) \in W$. Since $|s_j^r - r| \leq \varepsilon$, we have that $|r - t| < 2\varepsilon$. Likewise, $|s_k^r - r| \leq \varepsilon$, so $|t - s_k^r| < 3\varepsilon$. This implies that $\partial(f_k(t), f_k(s_k^r)) < 3\varepsilon$. Therefore $f_k(t)$ witnesses that $(f_j(s_j^r), f_k(s_k^r))$ is in A and therefore also V.

Since we can do this for any entourage W, we have that $\{f_i(s_i^r)\}_{i \in I}$ is a convergent net. $\square_{\text{claim}}$

Let $g(r)$ be the unique limit point of the net $\{f_i(s_i^r)\}_{i \in I}$ for each $r \in F$. By lower semi-continuity of ∂, $g(r)$ must be 1–Lipschitz, so it is an element of $\mathcal{I}(X)$ and the limit of the net $\{f_i\}_{i \in I}$.

Finally, it is immediate that any limit of a net of elements of $\mathcal{I}_r(X)$ is in $\mathcal{I}_r(X)$, so $\mathcal{I}_r(X)$ is closed for each r. $\square$

Corollary 4.4 *A net $\{f_i\}_{i \in I}$ of elements of $\mathcal{I}(X)$ converges if and only if for every entourage $V \subseteq X^2$ and $\varepsilon > 0$, there is an $i \in I$ such that for all $j, k \geq i$, $(f_j, f_k) \in U_{V,\varepsilon}$.*

Likewise, a filter $\mathcal{F}$ in $\mathcal{I}(X)$ converges to f if and only if for every entourage $V \subseteq X^2$ and $\varepsilon > 0$, the set $\{g \in \mathcal{I}(X) : (f, g) \in U_{V,\varepsilon}\}$ is in $\mathcal{F}$.

Proof This is immediate from Proposition 4.3 and the fact that entourages of the form $U_{V,\varepsilon} \subseteq \mathcal{I}(X)^2$ generate the uniform structure on $\mathcal{I}(X)$. $\square$

Proposition 4.5 *For any $r \geq 0$, ultrafilter $\mathcal{U}$ (with index set I), and family $\{f_i\}_{i \in I}$ of elements of $\mathcal{I}_r(X)$, the ultraproduct $\prod_{i \in I} \mathbb{I}(f_i)/\mathcal{U}$ is isomorphic to $\mathbb{I}(\lim_{i \to \mathcal{U}} f_i)$. In particular, for each $r \geq 0$, $\mathcal{I}_r(X)$ is compact.*

Proof Let M be the ultraproduct $\prod_{i\in I}\mathbb{I}(f_i)/\mathcal{U}$. For each $i\in I$, there is an isometric embedding of $\mathbb{I}(f_i)$ into $[0,r]$ defined by taking each $a\in\mathbb{I}(f_i)$ to $d(c^{\mathbb{I}(f_i)},a)$. This implies that the same is true of M, so we may identify the universe of M with a compact subset of $[0,r]$ containing $0=c^M$. Furthermore, for each $a\in M$, there is a unique element $f_{\mathcal{U}}(a)$ of X satisfying that $h(f_{\mathcal{U}}(a))=P^M_h(a)$ for each 1–Lipschitz $h:X\to\mathbb{R}$. $\mathrm{Th}(\mathcal{I}_r(X))$ ensures that $f_{\mathcal{U}}$ is a 1–Lipschitz function. It is clear by construction that M is isomorphic to $\mathbb{I}(f_{\mathcal{U}})$. Therefore we just need to show that $f_{\mathcal{U}}=\lim_{i\to\mathcal{U}}f_i$ (since the uniform structure on $\mathcal{I}_r(X)$ is Hausdorff, so ultrafilters have unique limits).

Fix an entourage $V\subseteq X^2$ and $\varepsilon>0$. We need to show that $\{i\in I:(f_i,f_{\mathcal{U}})\in U_{V,\varepsilon}\}$ is in $\mathcal{U}$. By [3, Theorem 1.6] and by shrinking ε if necessary, there is a finite set $\{h_j\}_{j<n}$ of 1–Lipschitz functions on X such that for any $p,q\in X$, if $|h_j(p)-h_j(q)|<\varepsilon$ for each $j<n$, then $(p,q)\in V$. Let $W=\{(p,q)\in X^2:(\forall j<n)|h_j(p)-h_j(q)|<\varepsilon\}$. We now have that $U_{W,\varepsilon}\subseteq U_{V,\varepsilon}$.

Assume for the sake of contradiction that $\{i\in I:(f_i,f_{\mathcal{U}})\in U_{W,\varepsilon}\}\notin\mathcal{U}$. This implies that for a $\mathcal{U}$–large set of $i\in I$, either

(1) there is an $s_i\in\mathrm{dom} f_i$ such that for all $t\in\mathrm{dom} f_{\mathcal{U}}$, either $(f_i(s_i),f_{\mathcal{U}}(t))\notin W$ or $|s_i-t|\geq\varepsilon$, or
(2) there is a $t_i\in\mathrm{dom} f_{\mathcal{U}}$ such that for all $s\in\mathrm{dom} f_i$, either $(f_i(s),f_{\mathcal{U}}(t_i))\notin W$ or $|s-t_i|\geq\varepsilon$.

Furthermore, one of these two cases must occur on a $\mathcal{U}$–large set of $i\in I$.

Assume that (1) occurs on a $\mathcal{U}$–large set of $i\in I$. Let $s_{\mathcal{U}}=\lim_{i\to\mathcal{U}}s_i$ (which we may regard both as an element of $[0,r]$ and M). Since $s_{\mathcal{U}}\in\mathrm{dom} f_{\mathcal{U}}$, we now have that for a $\mathcal{U}$–large set of $i\in I$, either $(f_i(s_i),f_{\mathcal{U}}(s_{\mathcal{U}}))\notin W$ or $|s_i-s_{\mathcal{U}}|\geq\varepsilon$. By the definition of W, if $(f_i(s_i),f_{\mathcal{U}}(s_{\mathcal{U}}))\notin W$, then there is a $j<n$ such that $|h_j(f_i(s_i))-h_j(f_{\mathcal{U}}(s_{\mathcal{U}}))|\geq\varepsilon$. Therefore either there is a $j<n$ such that on a $\mathcal{U}$–large set of $i\in I$, $|h_j(f_i(s_i))-h_j(f_{\mathcal{U}}(s_{\mathcal{U}}))|\geq\varepsilon$ or there is a $\mathcal{U}$–large set of $i\in I$ on which $|s_i-s_{\mathcal{U}}|\geq\varepsilon$. In any of these cases we get a contradiction.

Assume that (2) occurs on a $\mathcal{U}$–large set of $i\in I$. Fix a k in this set. Let $(u_i)_{i\in I}$ be a family (with $u_i\in\mathbb{I}(f_i)$ for each i) corresponding to $t_k\in M$ in the ultraproduct. We now have that on a $\mathcal{U}$–large set of $i\in I$, either $(f_i(u_i),f_{\mathcal{U}}(t_k))\notin W$ or $|u_i-t_k|\geq\varepsilon$. By the same argument as before we get a contradiction.

Therefore we must have that $\{i\in I:(f_i,f_{\mathcal{U}})\in U_{W,\varepsilon}\}\in\mathcal{U}$ and therefore $\{i\in I:(f_i,f_{\mathcal{U}})\in U_{V,\varepsilon}\}\in\mathcal{U}$. Since we can do this for any entourage V and $\varepsilon>0$, we have that $\lim_{i\to\mathcal{U}}f_i=f_{\mathcal{U}}$, as required.

Finally, since every ultrafilter on $\mathcal{I}_r(X)$ converges, it is compact. □

Proposition 4.5 implies that for each $r \geq 0$, $\mathcal{I}_r(X)$ can be topologically identified with $S_0(\mathrm{Th}(\mathcal{I}_r(X)))$ (the space of completions of $\mathrm{Th}(\mathcal{I}_r(X))$) via the map $f \mapsto \mathrm{Th}(\mathbb{I}(f))$.

5 The first-order theory of $\mathbb{F}(X)$

Given any $K, K' \in M \models \mathrm{Th}_{\mathbb{F}}(X)$, we will freely regard $[K, K']$ as an $\mathcal{L}_X(c)$–structure with the interpretation $c^{[K,K']} = K$. We say that $[K, K']$ and $[L, L']$ are *isomorphic* if they are isomorphic as $\mathcal{L}_X(c)$–structures.

Proposition 5.1 *Fix $M \models \mathrm{Th}_{\mathbb{F}}(X)$, $r > 0$, and $K, K' \in M$ with $d(K, K') < r$. $[K, K']$ is a model of $\mathrm{Th}(\mathcal{I}_r(X))$ (where we take the interpretation of c to be K in $[K, K']$).*

Proof Find an ultrafilter $\mathcal{U}$ (on an index set I) and an elementary embedding $f \colon M \preceq N := \mathbb{F}(X)^{\mathcal{U}}$. Identify M with its image under f. We can find families $\{K_i\}_{i\in I}$ and $\{K'_i\}_{i\in I}$ of elements corresponding respectively to K and K' satisfying $d(K_i, K'_i) < r$ for all $i \in I$. Since $\delta_{K_i,K'_i,r}(x)$ is the distance predicate of $[K_i, K'_i]^{\mathbb{F}(X)}$ for each i and likewise $\delta_{K,K',r}(x)$ is the distance predicate of $[K, K']^N$, we have that $[K, K']^N$ as an $\mathcal{L}_X(c)$–structure is isomorphic to the ultraproduct $\prod_{i\in I}[K_i, K'_i]^{\mathbb{F}(X)}/\mathcal{U}$, which is clearly a model of $\mathrm{Th}(\mathcal{I}_r(X))$. Finally, since $[K, K']^N$ is a compact set definable from K and K', we have that $[K, K']^N \subseteq M$, and we are done. □

Proposition 5.2 *Fix an ultrafilter $\mathcal{U}$ on an index set I and $K, K' \in M$ with $d(K, K') < \infty$. Fix families $\{K_i\}_{i\in I}$ and $\{K'_i\}_{i\in I}$ respectively corresponding to K and K' in the ultrapower. Fix a family $\{f_i\}_{i\in I}$ of elements of $\mathcal{I}(X)$ such that for a $\mathcal{U}$–large set of indices i, $[K_i, K'_i]$ is isomorphic to $\mathbb{I}(f_i)$.*

$\lim_{i\to\mathcal{U}} f_i$ exists, and $[K, K']$ is isomorphic as an $\mathcal{L}_X(c)$–structure to $\mathbb{I}(\lim_{i\to\mathcal{U}} f_i)$.

Proof Find $r > d(K, K')$. By the same argument as in the proof of Proposition 5.1, we have that $[K, K']$ is isomorphic to the ultraproduct $\prod_{i\in I}[K_i, K'_i]/\mathcal{U}$. The required conclusion now follows from Proposition 4.5. □

Now we come to the first point at which we actually need to assume that the metric ∂ on X is adequate.

Lemma 5.3 (Parallel itineraries) *(∂ adequate.) For any $f \in \mathcal{I}(X)$, any entourage V, and any $\varepsilon > 0$, there is an open neighborhood $O \ni f(0)$ such that for any $x \in O$, there is an itinerary $g \in \mathcal{I}(X)$ with $g(0) = x$ and $(f, g) \in U_{V,\varepsilon}$.*

Proof Recall that for any entourage $W \subseteq X^2$ and $z \in X$, we write $W(z)$ for $\{w : (w,z) \in W\}$. Note that if W is an open entourage (ie, an open neighborhood of the diagonal in X^2), then $W(z)$ is an open set.

Find an open entourage $W \subseteq V$ and a $\delta > 0$ with $\delta < \varepsilon$ small enough that for any $x \in X$, $\mathrm{cl}(W(x))^{\leq\delta} \subseteq V(x)$. (This is always possible by compactness.)

By compactness, we can find a finite set $F \subseteq \mathrm{dom} f$ such that $0 \in F$ and $\mathrm{dom} f \subseteq F^{<\frac{1}{2}\delta}$. Let $\{r_i\}_{i\leq n}$ be an increasing enumeration of F (with $r_0 = 0$). Fix $\gamma > 1$ small enough that $(\gamma - 1)\|f\| < \frac{1}{2}\delta$. (We will use γ at the end of the argument to ensure 1–Lipschitzness of the itineraries we construct to witness the desired property of $O \ni f(0)$.)

We need to build neighborhoods around each $f(r_i)$ small enough to witness the desired behavior of the neighborhood $O \ni f(0)$ that we will construct. Let $A_n = W(f(r_n))$. Then, for each $i < n$,

- if $f(r_i) = f(r_{i+1})$, let $A_i = W(f(r_i)) \cap A_{i+1}$ and
- if $f(r_i) \neq f(r_{i+1})$, let $A_i = W(f(r_i)) \cap A_{i+1}^{<\gamma\partial(f(r_i),f(r_{i+1}))}$.

Note that by adequacy of ∂, each A_i is an open neighborhood of $f(r_i)$.

Let $O = A_0$. For any $x \in O$, by construction, we can find a sequence $\{x_i\}_{i\leq n}$ such that

- $x_0 = x$,
- $x_i \in A_i \subseteq W(f(r_i))$ for each $i \leq n$, and
- $\partial(x_i, x_{i+1}) \leq \gamma\partial(f(r_i), f(r_{i+1}))$ for each $i < n$.

We are not quite done, as it may be the case that the function that maps r_i to x_i is not 1-Lipschitz. What we do have is that for each $i < n$, $\gamma^{-1}\partial(x_i, x_{i+1}) \leq \partial(f(r_i), f(r_{i+1})) \leq |r_{i+1} - r_i|$, so $\partial(x_i, x_{i+1}) \leq \gamma|r_{i+1} - r_i|$, which implies that the function that maps γr_i to x_i is 1–Lipschitz.

Let g be the element of $\mathcal{I}(X)$ with domain γF and with the property that for each $i \leq n$, $g(\gamma r_i) = x_i$. We want to show that $(g, f) \in U_{V,\varepsilon}$.

For each γr_i in $\mathrm{dom}\, g$, we have by construction that $|\gamma r_i - r_i| \leq (\gamma - 1)\|f\| < \frac{1}{2}\delta < \varepsilon$, and furthermore we have that $g(\gamma r_i) \in W(f(r_i)) \subseteq V(f(r_i))$.

For the other direction, we have that for any $s \in \mathrm{dom} f$, there is an $r_i \in F$ with $|s - r_i| < \frac{1}{2}\delta$. This implies that $|s - \gamma r_i| < \delta < \varepsilon$. By construction, we have that $f(r_i) \in W(g(\gamma r_i))$. Since $|s - r_i| < \frac{1}{2}\delta$, we have that $\partial(f(s), f(r_i)) < \frac{1}{2}\delta < \delta$ as well, so $f(s) \in (\mathrm{cl}\, W(g(\gamma r_i)))^{\leq\delta} \subseteq V(g(\gamma r_i))$, as required. Therefore $(g, f) \in U_{V,\varepsilon}$. □

Lemma 5.4 *For any $M \models \mathrm{Th}_{\mathbb{F}}(X)$ and any $a \in M$, there is a unique $p \in X$ with the property that for every 1–Lipschitz $f\colon X \to \mathbb{R}$, $f(p) = P_f^M(a)$.*

Proof By [3, Theorem 1.6], there is at most one such p. Find an ultrapower $\mathbb{F}^{\mathcal{U}}$ and an elementary embedding $g : M \preceq \mathbb{F}^{\mathcal{U}}$. Such a p exists for $g(a)$ by construction. By elementarity, the same p works for a as well. □

Definition 5.5 We write $\mathrm{tp}_X(a)$ for the point p shown to exist in Lemma 5.4.

Note that $\mathrm{tp}_X(a)$ is essentially the quantifier-free type of a. The notation is mostly to emphasize that we are thinking of it as an element of X.

Lemma 5.6 *(∂ adequate.) For any model $M \models \mathrm{Th}_{\mathbb{F}}(X)$, any $b \in M$, any $f \in \mathcal{I}(X)$ with $f(0) = b$, and any κ, there is an elementary extension $N \succeq M$ and a family $\{c_i\}_{i<\kappa}$ of elements of N such that*

- *for any $i < \kappa$, $[b, c_i]$ exists and is isomorphic to f and*
- *for any $i < j < \kappa$, $[b, c_i] \cap [b, c_j] = \{b\}$.*

Proof Clearly by compactness it is sufficient to show this with $\kappa = \aleph_0$. Fix $M \models \mathrm{Th}_{\mathbb{F}}(X)$, $b \in M$, and $f \in \mathcal{I}(X)$. Since b is an element of a model of $\mathrm{Th}_{\mathbb{F}}(X)$, there exists an ultrafilter $\mathcal{F}$ and an $a \in \mathbb{F}(X)^{\mathcal{F}}$ such that $a \equiv b$. Let a correspond to the family $\{K_i\}_{i\in I}$, where I is the index set of $\mathcal{F}$.

Let J be an index set large enough that we can find an enumeration $\{V_j, \varepsilon_j\}_{j\in J}$ of all pairs (V_j, ε_j) with $V_j \subseteq X^2$ an entourage and $\varepsilon_j > 0$. Lemma 5.3 implies that for each $j \in J$, we can find an open neighborhood $O \ni f(0)$ such that for any $x \in O$, there is an itinerary $g \in \mathcal{I}(X)$ such that $g(c) = x$ and $(f, g) \in U_{V_j,\varepsilon_j}$.

Fix some $j \in J$. By Lemma 5.3, we have that for an $\mathcal{F}$–large set of i, we can find a $g_{ij} \in \mathcal{I}(X)$ with $(f, g_{ij}) \in U_{V_j,\varepsilon_j}$ and a family $\{L^n_{ij}\}_{n<\omega} \subseteq \mathbb{F}(X)$ of extensions of K_i such that $\beta^{L^n_{ij}}(\|K_i\|) = n$ and $[K_i, L^n_{ij}]$ is isomorphic to $\mathbb{I}(g_{ij})$ for each $n < \omega$. In particular, note that for any $n < k < \omega$, $[K_i, L^n_{ij}] \cap [K_i, L^k_{ij}] = \{K_i\}$ and so $d(L^n_{ij}, L^k_{ij}) = d(L^n_{ij}, K_i) + d(K_i, L^k_{ij})$.

For each $n < \omega$, let L^n_j be the element of $\mathbb{F}(X)^{\mathcal{F}}$ corresponding to the family $\{L^n_{ij}\}_{i\in I}$. Let $g_j = \lim_{i\to\mathcal{F}} g_{ij}$. Note that by Proposition 4.5, $[K, L^n_j]$ is isomorphic to $\mathbb{I}(g_j)$ for each $n < \omega$ and that $(g_j, f) \in \mathrm{cl}(U_{V_j,\varepsilon_j})$. Furthermore, note that we have $d(L^n_j, L^k_j) = d(L^n_j, K) + d(K, L^k_j)$ and therefore $[K, L^n_j] \cap [K, L^k_j] = \{K\}$ for each $n < k < \omega$.

Let $\mathcal{G}$ be an ultrafilter on J satisfying that for each $j \in J$, the set $\{j' \in J : U_{V_{j'},\varepsilon_{j'}} \subseteq U_{V_j,\varepsilon_j}\}$ is in $\mathcal{G}$. Note that $\lim_{j\to\mathcal{G}} g^n_j = f$ for all $n < \omega$.

Let N be the ultrapower $(\mathbb{F}(X)^{\mathcal{F}})^{\mathcal{G}}$. Identify M with its image in N under the diagonal embedding, so that we can regard K as an element of N in the obvious way. For each $n < \omega$, let L^n be the element of N corresponding to the family $\{L^n_j\}_{j\in J}$. We clearly have that for any $n < k < \omega$, $d(L^n, L^k) = d(L^n, K) + d(K, L^k)$, so $[K, L^n] \cap [K, L^k] = \{K\}$. By Proposition 4.5, we have that each $[K, L^n]$ is isomorphic to $\mathbb{I}(\lim_{j\to\mathcal{G}} g^n_j) = \mathbb{I}(f)$, so $\{L^n\}_{n<\omega}$ is the required family of elements. □

Lemma 5.7 *(∂ adequate.) Let $\mathbb{M}$ be a $(2^{\aleph_0+|X|})^+$–saturated, $(2^{\aleph_0+|X|})^+$–homogeneous monster model of $\mathrm{Th}_{\mathbb{F}}(X)$. Let $\bar{a} = \bar{a}^0\bar{a}^1 \dots \bar{a}^{n-1}$ and $\bar{b} = \bar{b}^0\bar{b}^1 \dots \bar{b}^{n-1}$ be finite tuples of elements of $\mathbb{M}$ partitioned into finite distance classes. Let $A = \mathrm{ccl}(\bar{a}^0) \cup \mathrm{ccl}(\bar{a}^1) \cup \dots \cup \mathrm{ccl}(\bar{a}^{n-1})$ and $B = \mathrm{ccl}(\bar{b}^0) \cup \mathrm{ccl}(\bar{b}^1) \cup \dots \cup \mathrm{ccl}(\bar{b}^{n-1})$. Assume that there is an $\mathcal{L}_X$–isomorphism $f\colon A \cong B$ such that for each i, j, $f(a^j_i) = b^j_i$. Then for any $c \in \mathbb{M}$, there exists an $e \in \mathbb{M}$ such that*

- *if c is not in the finite distance class of any element of $\bar{a}$, then e is not in the finite distance class of any element of $\bar{b}$, and the map $g\colon Ac \to Be$ extending f by letting $g(c) = e$ is an $\mathcal{L}_X$–isomorphism and*
- *if c is in the finite distance class of $\bar{a}^i$, then e is in the finite distance class of $\bar{b}^i$, and there is an $\mathcal{L}_X$–isomorphism $g\colon A \cup \mathrm{ccl}(\bar{a}^i c) \cong B \cup \mathrm{ccl}(\bar{b}^i e)$ extending f such that $g(c) = e$.*

Proof If c is not in the same finite distance class as any element of $\bar{a}$, then we can easily find $e \in \mathbb{M}$ not in the same finite distance class as any element of $\bar{b}$ such that $\mathrm{tp}_X(e) = \mathrm{tp}_X(c)$. Then g extending f to Ac in the obvious way is clearly an $\mathcal{L}_X$–isomorphism.

If c is in the same finite distance class as $\bar{a}^i$, then by Proposition 3.16, there is a unique element $c' \in \mathrm{ccl}(\bar{a}^i)$ with $d(c, c') = d(c, \mathrm{ccl}(\bar{a}^i))$. Let h be the element of $\mathcal{I}(X)$ corresponding to $[c', c]$. By assumption, we have that $e' \coloneqq f(c') \in B$ has $\mathrm{tp}_X(e') = \mathrm{tp}_X(c')$, so by Lemma 5.6, there is a family $\{e_i\}_{i<(2^{\aleph_0})^+}$ of elements of $\mathbb{M}$ such that for each $i < \left(2^{\aleph_0}\right)^+$, $[e', e_i]$ corresponds to h in $\mathcal{I}(X)$ and for each $i < j < \left(2^{\aleph_0}\right)^+$, $[e', e_i] \cap [e', e_j] = \{e'\}$. Since the cardinality of $\mathrm{ccl}(\bar{b}^i)$ is at most $2^{\aleph_0}$, by the pigeonhole principle, there must be some $i < \left(2^{\aleph_0}\right)^+$ such that $\mathrm{ccl}(\bar{b}^i) \cap [e', e_i] = \{e'\}$. Let e be that e_i.

We can extend f to g by setting $g(x)$, for each $x \in [c', c] \setminus \{c'\}$, to the unique element y of $[e', e] \setminus \{e'\}$ such that $d(c', x) = d(e', y)$. Since $[c', c]$ and $[e', e]$ both correspond to h in $\mathcal{I}(X)$, we have that g is an $\mathcal{L}_X$–isomorphism. Finally, we clearly have that $g(c) = e$. □

Proposition 5.8 *(∂ adequate.) For any finite tuple $\bar{a}$ in any model $M \models \mathrm{Th}_{\mathbb{F}}(X)$, $\mathrm{tp}(\bar{a})$ is uniquely determined by the partitioning of $\bar{a}$ into finite distance classes and the $\mathcal{L}_X(\bar{b})$–isomorphism type of each $\mathrm{ccl}(\bar{b})$ for $\bar{b}$, a finite distance class of $\bar{a}$.*

Proof This follows from Lemma 5.7 and a back-and-forth argument. □

Corollary 5.9 *(∂ adequate.) For any $a \in M \models \mathrm{Th}_{\mathbb{F}}(X)$, $\mathrm{tp}(a)$ is uniquely determined by $\mathrm{tp}_X(a)$.*

Proof Clearly we have that if $\mathrm{tp}_X(a) \neq \mathrm{tp}_X(b)$, then $\mathrm{tp}(a) \neq \mathrm{tp}(b)$. Conversely, if $\mathrm{tp}_X(a) = \mathrm{tp}_X(b)$, then by Proposition 5.8, we have that $\mathrm{tp}(a) = \mathrm{tp}(b)$. □

Corollary 5.10 *(∂ adequate.) For any $n < \omega$, if $p(\bar{x})$ is an n–type of $\mathrm{Th}_{\mathbb{F}}(X)$ over $\varnothing$ with pairwise finite distances, then $p(\bar{x})$ is realized in $\mathbb{F}(X)$.*

Proof First note that for 1–types, this follows immediately from Corollary 5.9. Furthermore, we may assume that the $K \in \mathbb{F}(X)$ realizing p restricted to the first variable x_0 satisfies $\|K\| = 0$.

Assume that we have shown the statement for n and let $p(x_0, x_1, \dots, x_n)$ be an $n+1$–type. Let $\bar{a}$ be a realization of p in the monster. By the induction hypothesis, we can find $K_0, K_1, \dots, K_{n-1}$ such that $\|K_0\| = 0$ and $K_0 \dots K_{n-1} \models p(x_0, \dots, x_{n-1})$. Let b be the element of $\mathrm{ccl}(a_0 \dots a_{n-1})$ closest to a_n (shown to exist in Proposition 3.16). Since $K_0 \dots K_{n-1}$ realizes the same type as $a_0 \dots a_{n-1}$, there is an $\mathcal{L}_X$–isomorphism $f : \mathrm{ccl}(a_0 \dots a_{n-1}) \to \mathrm{ccl}(K_0 \dots K_{n-1})$ with $f(a_i) = K_i$ for each $i < n$. Let $L = f(b)$. If $a_n = b$, then we are done. Otherwise, we can find a K_n extending L satisfying that $[L, K_n]$ is isomorphic to $[b, a_n]$. Furthermore, by letting $\beta^{K_n}(\|L\|)$ be sufficiently large, we can ensure that $\mathrm{ccl}(K_0 \dots K_{n-1}) \cap [L, K_n] = \{L\}$. Therefore f extends to an $\mathcal{L}_X$–isomorphism $g : \mathrm{ccl}(a_0 \dots a_n) \to \mathrm{ccl}(K_0 \dots K_n)$ satisfying $g(a_n) = K_n$. By Proposition 5.8, we have that $K_0 \dots K_n$ realizes the same type as $a_0 \dots a_n$, namely $p(\bar{x})$. □

Note that not all n–types may be realized in $\mathbb{F}(X)$. For instance, if X is a single point, then $\mathbb{F}(X)$ has a single finite-distance component. This could be fixed by duplicating the roots (ie, elements of length 0) in the construction of $\mathbb{F}(X)$, but we only need Corollary 5.10 for types with pairwise finite distances.

6 Stability of $\mathrm{Th}_{\mathbb{F}}(X)$

From now on we will assume that ∂ is an adequate metric.

Lemma 6.1 *For any model $M \models \mathrm{Th}_{\mathbb{F}}(X)$, any elementary extension $N \succeq M$, and any $a \in N$, either a does not have finite distance to any element of M or there is a unique $e \in M$ with minimal distance to a.*

Proof Consider the set $F \coloneqq \bigcap\{[a,c]^N : c \in M, d(a,c) < \infty\}$. Since this is the intersection of a family of compact sets, it itself is compact. Since F is compact, it contains an element e such that $d(e,M)$ is minimized.

Claim. $d(e,M) = 0$, or, in other words, $e \in M$.

Proof of claim. Suppose that $d(e,M) > 0$. Since $x \mapsto d(x,M)$ is a continuous function, there must be by compactness some finite set $M_0 \subset M$ of elements with finite distance to a such that $F_0 \coloneqq \inf\{f \in \bigcap\{[a,c]^N : c \in M_0\}\} > 0$. By Proposition 3.16, there is a unique element $g \in \mathrm{ccl}(M_0) \subset M$ of minimal distance to a. It must be the case that $g \notin F_0$, so there must be some $m \in M_0$ such that $g \notin [m,a]$. Let $h = Y(a,g,m)$.

It is easy to check that in $\mathbb{F}(X)$, for any A, B, and C with pairwise finite distances, $Y(A,B,C) \in [A,C]$. By Corollary 5.10, all 3–types with pairwise finite distances are realized in $\mathbb{F}(X)$, so we have that $Y(A,B,C) \in [A,C]$ holds for all models of $\mathrm{Th}_{\mathbb{F}}(X)$ as well. Therefore we have that $h \in [m,g]$. Since $[m,g]$ is contained in the algebraic closure of mg,[7] we have that $[m,g] \subset M$. By Lemma 3.14, we have that $[m,g] \subseteq \mathrm{ccl}(M_0)$, and so $h \in \mathrm{ccl}(M_0)$, but $d(h,a) < d(g,a)$, which is a contradiction. $\square_{\text{claim}}$

Claim. e is unique.

Proof of claim. Suppose that there are distinct e and e' in $F \cap M$. This implies that for any $m \in M$ with $d(m,a) < \infty$, e and e' are both in $[m,a]$. One of e and e' must be closer to a. Assume without loss of generality that $d(e,a) < d(e',a)$. Then we have that $e' \notin [e,a] \subseteq F$, which is a contradiction. $\square_{\text{claim}}$

The argument for the last claim also establishes that for any $m \in M$, $d(m,a) \geq d(e,a)$. $\square$

Lemma 6.2 *Fix $a,b,c,e \in M \models \mathrm{Th}_{\mathbb{F}}(X)$ with pairwise finite distance. Let $c' = Y(a,b,c)$, and let $e' = Y(a,b,e)$. If $c' \neq e'$, then $d(c,e) = d(c,c') + d(c',e') + d(e',e)$.*

[7] $[m,g]$ is actually identical to $\mathrm{acl}(mg)$ and $\mathrm{dcl}(mg)$, but we have not established this.

Proof It is easy to verify that this is true in $\mathbb{F}(X)$. The full statement thereby follows from Corollary 5.10. □

Lemma 6.3 *Any $\mathbb{R}$–tree embeds isometrically into a model of $\mathrm{Th}_{\mathbb{F}}(X)$.*

Proof Fix $p \in X$. For each $r \geq 0$, let $Y_r \subseteq \mathbb{F}(X)$ be the set of all elements K with $\operatorname{dom} K = [0, s]$ for some $s \leq r$ and $\chi^K(t) = p$ for all $t \in \operatorname{dom} K$. It is straightforward to show that Y_r (with the unique element $K_0 \in Y_r$ satisfying $\|K_0\| = 0$ chosen as a root) is a richly branching $\mathbb{R}$–tree of height r in the sense of [7, Definition 7.1]. Therefore $(Y_r, K_0, d_{2r}) \models \mathrm{rb}\mathbb{R}\mathrm{T}_r$ and so by [7, Theorem 7.12], every pointed $\mathbb{R}$–tree of radius at most r isometrically embeds into a model of $\mathrm{Th}_{\mathbb{F}}(X)$. By compactness, this implies that every $\mathbb{R}$–tree embeds isometrically into a model of $\mathrm{Th}_{\mathbb{F}}(X)$. □

Lemma 6.4 *For any $\kappa \geq |\mathcal{L}_X|$, there is a model M of $\mathrm{Th}_{\mathbb{F}}(X)$ with density character κ such that $|M| = \kappa^{\aleph_0}$.*

Proof By the argument in the proof of Theorem 8.10 in [7], there is an $\mathbb{R}$–tree N with density character κ and cardinality $\kappa^{\aleph_0}$. By Lemma 6.3, we can find $M \models \mathrm{Th}_{\mathbb{F}}(X)$ and an isometric embedding $f\colon N \to M'$. We can apply Löwenheim–Skolem to find $M' \preceq M$ with $f(N) \subseteq M'$ such that the density character of M' is κ. Any metric space with density character κ has cardinality at most $\kappa^{\aleph_0}$, so we are done. □

Proposition 6.5 *For any model $M \models \mathrm{Th}_{\mathbb{F}}(X)$, the elements of $S_1(M)$ are precisely*

- *the realized types in M,*
- *types $p_{m,f}$ for each pair $m \in M$ and $f \in \mathcal{I}(X)$ with $f(0) = \mathrm{tp}_X(m)$ and $\|f\| > 0$, and*
- *types q_x for each $x \in X$,*

where

- *$p_{m,f}$ is the type of an element $a \in N \succ M$ whose unique nearest element in M is m and which satisfies that $[m, a]$ is isomorphic to $\mathbb{I}(f)$ (with f in $\mathcal{I}(X)$) and*
- *q_x is the type of an element $b \in N \succ M$ with $\mathrm{tp}_X(b) = x$ and $d(b, M) = \infty$.*

Furthermore, the metric[8] δ on $S_1(M)$ is given by

- *$\delta(p_{m,f}, p_{m',f'}) = \min\{\|f\| + d(m, m') + \|f'\|, \operatorname{diam} X\}$ if $m \neq m'$,*
- *$\delta(p_{m,f}, p_{m,f'}) = \min\{\|f \sqcap f'\|, \operatorname{diam} X\}$, where $f \sqcap f'$ is the longest common initial segment of f and f',*
- *$\delta(q_x, q_{x'}) = \partial(x, x')$, and*

[8] Recall that the 'official' metric in $\mathcal{L}_X$ is $d_{\operatorname{diam} X}$.

- $\delta(m, q_x) = \delta(p_{m,f}, q_x) = \operatorname{diam} X$

for any $m, m' \in M$, $x, x' \in X$, and $f, f' \in \mathcal{I}(X)$. So in particular,

$$|M| + \#^{\mathrm{dc}} X \leq \#^{\mathrm{dc}} S_1(M) \leq |\mathcal{I}(X)| \cdot |M| + \#^{\mathrm{dc}} X$$

and $\mathrm{Th}_{\mathbb{F}}(X)$ is strictly stable (where $\#^{\mathrm{dc}} Y$ is the metric density character of Y).

Proof Clearly every type $p(x)$ in $S_1(M)$ is either realized, satisfies $d(x, m) < \infty$ for some $m \in M$, or satisfies $d(x, m) = \infty$ for every $m \in M$.

First assume that $p(x)$ and $p'(x)$ are not realized but satisfies $d(x, m) < \infty$ for some $m \in M$. Furthermore let $a \models p$ and $a' \models p'$. Assume that the unique nearest element b of M to a is also the unique nearest element to a'. (Recall that b exists by Lemma 6.1). Finally, assume that $[b, a]$ and $[b, a']$ are isomorphic. We need to show that $p = p'$. By Proposition 5.8, we have that for any finite $M_0 \subseteq M$, there is an automorphism of the monster taking $M_0 a$ to $M_0 a'$. Therefore $a \equiv_M a'$, and so $p = p'$.

Now assume that $p(x)$ and $p'(x)$ are not realized and satisfy $d(x, m) = \infty$ for every $m \in M$ (ie, $d_r(x, m) = r$ for every r and $m \in M$). Let $a \models p$ and $a' \models p'$. Assume that $\mathrm{tp}_X(a) = \mathrm{tp}_X(a')$. Again, we have by Proposition 5.8 that for every finite $M_0 \subseteq M$, there is an automorphism of the monster taking $M_0 a$ to $M_0 a'$. Therefore $a \equiv_M a'$, and so $p = p'$.

For the metric δ on $S_1(M)$, the last three bullet points are clearly correct. The first bullet point is clearly an upper bound, so we just need to show that a smaller distance cannot be achieved. If $a, b \in N \succ M$ have nearest points $c, e \in M$, respectively, then these are also their nearest points on $[c, e] \subset M$, so by Lemma 6.2, $d(a, b) = d(a, c) + d(c, e) + d(e, b)$, as required.

The bounds on the density character of $S_1(M)$ are obvious. In particular, since $|M| \leq (\#^{\mathrm{dc}} M)^{\aleph_0}$, we have that $\mathrm{Th}_{\mathbb{F}}(X)$ is stable. To see that $\mathrm{Th}_{\mathbb{F}}(X)$ is strictly stable, we have by Lemma 6.4 that for each $\kappa \geq |\mathcal{L}_X|$, there is an $M \models \mathrm{Th}_{\mathbb{F}}(X)$ with $\#^{\mathrm{dc}} M = \kappa$ and $|M| = \kappa^{\aleph_0}$. Therefore

$$\#^{\mathrm{dc}} S_1(M) \geq |M| + \#^{\mathrm{dc}} X = \kappa^{\aleph_0} + \#^{\mathrm{dc}} X$$

and so $\mathrm{Th}_{\mathbb{F}}(X)$ is strictly stable. □

7 Main theorem

Theorem 7.1 *For any compact topometric space (X, τ, ∂) with an adequate metric, there is a stable continuous first-order theory T such that $S_1(T)$ is topometrically isomorphic to (X, τ, ∂).*

Proof By Lemma 0.5, X has finite diameter, so we can form the theory $\mathrm{Th}_{\mathbb{F}}(X)$. There is clearly a continuous 1–Lipschitz map f from $S_1(T)$ to X. We have by Corollary 5.9 that f is a bijection, so it is a topological isomorphism. For any $a, b \in X$, there are, by construction, K and L in $\mathbb{F}(X)$ with $\mathrm{tp}_X(K) = a$ and $\mathrm{tp}_X(L) = b$ such that $d_{\mathrm{diam}\,X}(K, L) = d(K, L) = \partial(a, b)$. Therefore f is an isometry as well. Finally, by Proposition 6.5, $\mathrm{Th}_{\mathbb{F}}(X)$ is stable. □

It is natural to wonder if our main theorem can be improved by constructing a superstable theory T with $S_1(T)$ isomorphic to a given X. In other words, are there any non-trivial restrictions on the topometry type of $S_1(T)$ for superstable T? Clearly if X is not CB-analyzable,[9] then any such T cannot be ω–stable or totally transcendental, but it is also possible that this is the only obstruction.

Question 7.2 *If (X, τ, ∂) is a compact topometric space with an adequate metric, is there a superstable theory T such that $S_1(T)$ is isomorphic to X?*

If (X, τ, ∂) is CB-analyzable, is there a totally transcendental theory T such that $S_1(T)$ is isomorphic to X?

For comparison, note that every totally disconnected compact Hausdorff space is $S_1(T)$ for a superstable theory T, and every scattered compact Hausdorff space is $S_1(T)$ for a totally transcendental theory T.

There is also the task of characterizing higher type spaces. Even for 2–types, there are new restrictions on what topometry types are possible. If T has models with more than one element, then $S_2(T)$ has a non-trivial definable set, namely, $d(x, y) = 0$.

References

[1] **I Ben Yaacov**, *On perturbations of continuous structures*, Journal of Mathematical Logic 08 (2008) 225–249; https://doi.org/10.1142/S0219061308000762

[2] **I Ben Yaacov**, *Topometric spaces and perturbations of metric structures*, Logic and Analysis 1 (2008) 235–272; http://doi.org/10.1007/s11813-008-0009-x

[3] **I Ben Yaacov**, *Lipschitz functions on topometric spaces*, Journal of Logic and Analysis 5 (2013), paper 8; http://doi.org/10.4115/jla.2013.5.8

[4] **I Ben Yaacov**, *Fraïssé limits of metric structures*, The Journal of Symbolic Logic 80 (2015) 100–115; http://doi.org/10.1017/jsl.2014.71

[9]This is the topometric generalization of scatteredness. See [2].

[5] **I Ben Yaacov**, **A Berenstein**, **C W Henson**, **A Usvyatsov**, *Model theory for metric structures*, from: "Model Theory with Applications to Algebra and Analysis", (Z Chatzidakis, D Macpherson, A Pillay, A Wilkie, editors), London Mathematical Society Lecture Note Series 2, Cambridge University Press (2008) 315–427; http://doi.org/10.1017/CBO9780511735219.011

[6] **I Ben Yaacov**, **J Melleray**, *Grey subsets of Polish spaces*, The Journal of Symbolic Logic 80 (2015) 1379–1397; http://doi.org/10.1017/jsl.2014.60

[7] **S Carlisle**, **C W Henson**, *Model theory of* $\mathbb{R}$*-trees*, Journal of Logic and Analysis 12 (2020); http://doi.org/10.4115/jla.2020.12.3

[8] **I Chiswell**, *Introduction to* Λ*-Trees*, World Scientific (2001); http://doi.org/10.1142/4495

[9] **J Hanson**, *Strongly Minimal Sets and Categoricity in Continuous Logic*, arXiv e-prints (2020) arXiv:2011.00610, accepted at *Memoirs of the AMS*

[10] **J Roe**, *Lectures on Coarse Geometry*, American Mathematical Society (2003); http://doi.org/10.1090/ulect/031

Department of Mathematics, Iowa State University, 396 Carver Hall, 411 Morrill Road, Ames, IA 50011, USA

jameseh@iastate.edu

Received: 24 September 2021 Revised: 2 January 2025