TOPOMETRIC CHARACTERIZATION OF TYPE SPACES IN CONTINUOUS LOGIC

JAMES HANSON

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

Introduction

Continuous first-order logic, introduced in [4], is a generalization of discrete first-order logic suited for studying structures with natural metrics, such as C*-algebras, valued fields, and $\mathbb{R}$-trees. 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.

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 $d(x, D) := \inf_{y\in D} d(x, y)$ is a definable predicate.

In discrete logic, definable sets have a purely topological characterization in terms of clopen subsets of type space. In continuous logic, there is important extra structure on type spaces, namely the induced metric given by

$$d(p, q) := \inf\{d(a, b) : a \models p,\, b \models q\}.$$

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 abstracted to 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 (i.e., has $\{(x, y) \in X^2 : \partial(x, y) \leq \varepsilon\}$ closed for every $\varepsilon > 0$).

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

Date: June 22, 2021.

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

In [6], the author showed that in ω-stable (and more generally totally transcendental) theories, definable sets are prevalent in a precise sense (referred to there as *dictionaricness*) that permits, at least approximately, many of the manipulations of definable sets that are trivial in discrete logic. In general, though, it is known that there is a weakly minimal theory T such that $S_1(T)$ has cardinality $2^{\aleph_0}$ but no non-trivial definable sets. (See Example 4.2 in [6].)

These facts motivate a desire to solve the following question: Which topometric spaces are type spaces of some continuous theory, and what restrictions, if any, do various model-theoretic tameness conditions impose on the topometry types of type spaces? Despite the existence of superstable theories that are poorly behaved with regards to definable sets, there might be, in principle, some subtler regularity imposed by stability or superstability on them.

In discrete logic, this question is very easy to answer. Given 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 (i.e., has ordinal Cantor-Bendixson rank), then T can be taken to be totally transcendental.

The metric d on a given type space $S_n(T)$ has an additional regularity property not shared by compact topometric spaces in general, referred to provisionally in [1] as openness.

Definition 0.2. A topometric space (X, τ, ∂) has an *open* 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 openness of metric precisely characterize the topometry types of type spaces in continuous logic. Furthermore, we will show that any compact topometric space with an open 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 topometry types of type spaces in totally transcendental and superstable theories.

1. Generic Forest Structures

Definition 1.1. 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 U_f for each continuous function $f : X \to \mathbb{R}$, 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 continuous $f : X \to \mathbb{R}$, we pick a non-decreasing modulus of uniform continuity $\alpha_{U_f} : [0, \operatorname{diam} X] \to \mathbb{R}$ such that $f : X \to \mathbb{R}$ is α_{U_f}-uniformly continuous. For any L-Lipschitz function f, we will require that $\alpha_{U_f}(x) = Lx$.[2] The 'official' metric of $\mathcal{L}_X$ is $d_{\operatorname{diam} X}$.

We define the *generic forest structure*[3] *of* X, written $\mathfrak{F}(X)$, as the $\mathcal{L}_X$-structure whose universe consists of triples $K = (\pi(K), K_X, K_\omega)$ where[4]

- $\pi(K)$ is a compact subset of $\mathbb{R}_{\geq 0}$ containing 0,
- $K_X : \pi(K) \to X$ is a 1-Lipschitz function, and

[2] Recall that continuous functions on compact topometric spaces are automatically uniformly continuous with regards to the metric.

[3] The motivation for this name is that these structures are always embeddable in an $\mathbb{R}$-forest. Strictly speaking the name should be something like *generic* $\mathbb{R}$*-forest-embeddable structure*, but we have decided that this particular pedantry is not worthwhile.

[4] The use of ω in this definition is arbitrary. Any infinite cardinal would be sufficient.

- $K_\omega : (\pi(K) \setminus \{\sup \pi(K)\}) \to \omega$ is an arbitrary function.

For K and L in $\mathfrak{F}(X)$, we say that L *extends* K, written $K \sqsupseteq L$, if $\pi(K) \subseteq \pi(L)$, $L_X \upharpoonright \pi(K) = K_X$, and $L_\omega \upharpoonright (\pi(K) \setminus \{\sup \pi(K)\}) = K_\omega$.

We will write $|K|$ for $\sup \pi(K)$, which we will call the *length* of K.

We say that two elements K and K' of $\mathfrak{F}(X)$ are in the same *finite distance component* of $\mathfrak{F}(X)$ (or that they have *finite distance*) if $K_X(0) = K'_X(0)$.

For any $\mathcal{L}_X$-predicate U_f, we set

$$U_f^{\mathfrak{F}(X)}(K) = f(K_X(|K|)).$$

For any $K \in \mathfrak{F}(X)$ and $r \in [0, |K|]$, we write $K \upharpoonright [0, r]$ for the unique maximal element of $\mathfrak{F}(X)$ such that $L \sqsupseteq K$ and $|L| \leq r$. We call elements of this form *initial segments* of K.

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 $r \in \pi(K) \cap \pi(K')$ and $K \upharpoonright [0, r] = K' \upharpoonright [0, r]$. The element $K \upharpoonright r = K' \upharpoonright r$ is the *longest common initial segment of* K *and* K', written $K \sqcap K'$ if it exists.

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

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

Proposition 1.2. *In any* $\mathfrak{F}(X)$, 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 $K(0) = K'(0) = K''(0)$. Without loss of generality, assume that $|K \sqcap K'| \leq |K' \sqcap K''|$. We 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 $||K| - |K'|| \leq d(K, K')$.

Lemma 1.3. *For any set* $\mathbf{K}$ *of elements of* $\mathfrak{F}(X)$ *with pairwise finite distance, there is a unique longest common initial segment* $\bigsqcap \mathbf{K}$ *of* $\mathbf{K}$.

Proof. Let

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

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

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

Lemma 1.4. *For any set* $\mathbf{K}$ *of elements of* $\mathfrak{F}(X)$ *with pairwise finite distance,*

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

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

$$|K| + |K'| - 2|K \sqcap K'| \leq \operatorname{diam}\mathbf{K},$$
$$\frac{1}{2}(|K| + |K'| - \operatorname{diam}\mathbf{K}) \leq |K \sqcap K'|.$$

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

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

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

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

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

$$\left|\bigsqcap \mathbf{K}\right| \geq \frac{1}{2}\left(2\sup_{K\in\mathbf{K}} - 2\operatorname{diam}\mathbf{K}\right) = \sup_{K\in\mathbf{K}} |K| - \operatorname{diam}\mathbf{K}.$$ ☐

Corollary 1.5. *If* $\mathbf{K} \subseteq \mathfrak{F}(X)$ *has diameter at most* $r < \infty$*, then for any* $K \in \mathbf{K}$*,* $d(K, \bigsqcap \mathbf{K}) \leq \operatorname{diam}\mathbf{K}$.

Proof. For any $K \in \mathbf{K}$, we have that

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

as required. ☐

Proposition 1.6. *The metric* d *on* $\mathfrak{F}(X)$ *is complete.*

Proof. Let $\{K_i\}_{i<\omega}$ be a Cauchy sequence in $\mathfrak{F}(X)$. 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 1.5, we have that $d(K_i, L_i) \to 0$ as $i \to \infty$.

Let $A = \bigcup_{i<\omega} L_i$. Either $|A| = 0$, in which case do nothing, or $|A| > 0$, in which case there is a unique pair $(|A|, x)$ with $x \in X$ which makes $B = A \cup \{(|A|, x)\}$ an element of $\mathfrak{F}(X)$. By construction we have that B is the limit of the sequence $\{L_i\}_{i<\omega}$ and therefore of the sequence $\{K_i\}_{i<\omega}$ as well. ☐

Proposition 1.7. *For any continuous function* $f : X \to \mathbb{R}$*, the interpretation* $U_f^{\mathfrak{F}(X)}$ *is* α_{U_f}*-uniformly continuous.*

Proof. Recall that we have chosen α_{U_f} so that f is α_{U_f}-uniformly continuous on X. Also, note that we are really talking about uniform continuity with regards to the 'official' metric $d_{\operatorname{diam} X}$.

Fix $K, K' \in \mathfrak{F}(X)$. 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 K_X and K'_X are 1-Lipschitz, we have that

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

Therefore

$$\begin{aligned}|U_f(K) - U_f(K')| &\leq \alpha_{U_f}(\partial(K_X(|K|), K'_X(|K|))) \\ &\leq \alpha_{U_f}(d(K, K')).\end{aligned}$$

So since α_{U_f} is non-decreasing, U_f is α_{U_f}-uniformly continuous in $\mathfrak{F}(X)$. □

Corollary 1.8. *$\mathfrak{F}(X)$ is an $\mathcal{L}_X$-structure.*

Proof. Given Proposition 1.7, 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 1.9. For any X, let $T(X)$ be $\operatorname{Th}(\mathfrak{F}(X))$.

2. The Space of Paths

In this section we develop some machinery needed to analyze arbitrary models of $T(X)$.

Definition 2.1. For any compact topometric space (X, τ, ∂), let the *set of paths in* X, written $\mathfrak{P}(X)$, be the collection of all partial 1-Lipschitz functions $f : \subseteq \mathbb{R}_{\geq 0} \to X$ with compact domain containing 0. We write $|f|$ for $\sup \operatorname{dom} f$.

Note that for any $K \in \mathfrak{F}(X)$, we have that K_X is an element of $\mathfrak{P}(X)$. Also note that elements of $\mathfrak{P}(X)$ are automatically *topologically* continuous as well, since the ordinary topology is coarser than the metric topology.

We put a uniform structure (and therefore a topology) on $\mathfrak{P}(X)$ generated by entourages of the form $U_{V,\varepsilon}$ for V an entourage in X^2 and $\varepsilon > 0$, where $(f, g) \in U_{V,\varepsilon}$ 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 $\mathfrak{P}(X)$, note that

- $U_{V \cap W, \min\{\varepsilon, \delta\}} \subseteq U_{V,\varepsilon} \cap U_{W,\delta}$ and

- if $W^{\circ 2} \coloneqq \{(x,z) : (\exists y \in X)(x,y) \in W \wedge (y,z) \in W\} \subseteq V$, then $U_{W,\varepsilon/2}^{\circ 2} \subseteq U_{V,\varepsilon}$.

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, 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 entourage V, we write $V(x)$ for the set $\{y \in X : (x,y) \in V\}$.

Proposition 2.2. *The uniform structure on $\mathfrak{P}(X)$ is Hausdorff and complete.*

Proof. First to see that the uniform structure on $\mathfrak{P}(X)$ is Hausdorff, let f and g be distinct elements of $\mathfrak{P}(X)$. If $\operatorname{dom} f \neq \operatorname{dom} g$, then there must be an $\varepsilon > 0$ small enough that $(f,g) \notin U_{V,\varepsilon}$ for any entourage V. 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 $(f,g) \in U_{V,\varepsilon}$. By definition, this means that there is some $s \in \operatorname{dom} g$ with $|r - s| < \varepsilon$ such that $(f(r), g(s)) \in V$, i.e., $g(s) \in V(f(r))$. 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 $(f,g) \notin U_{V,\varepsilon}$.

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

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 $r \in F$, define a net $\{x_i^r\}_{i \in I}$ of points in X by setting x_i^r to $f_i(s_i^r)$ where s_i^r is 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.)

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

Proof of claim. Fix an entourage $V \subseteq X^2$. Find an entourage W 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 $j, k \geq i$, $f_j \in U_{f_i,W,\varepsilon}$ 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 $(x_j^r, f_k(t)) \in W$. 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 $\partial(f_k(t), x_k^r) < 3\varepsilon$. Therefore $f_k(t)$ witnesses that (x_j^r, x_k^r) is in A and therefore also V.

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

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

Note that the converse of Proposition 2.2 is immediate, so we have a characterization of topological convergence in $\mathfrak{P}(X)$.

Proposition 2.3. *For any $r \geq 0$, there is a theory T_r in the language $\mathcal{L}_X(c)$ (i.e., $\mathcal{L}_X$ with a single constant added) such that the models of T_r can be naturally* *identified with the elements of* $\mathfrak{P}_r(X) \coloneqq \{f \in \mathfrak{P}(X) : |f| \leq r\}$ *such that the topology on* $\mathfrak{P}_r(X)$ *agrees with the topology induced by ultraproducts. In particular, each* $\mathfrak{P}_r(X)$ *is compact.*

Proof. Consider the theory T_r containing $d_{r+1}(x,y) \leq r$ and

$$d_r(x,y) + d_r(x,z) + d_r(y,z) = 2\max\{d_r(x,y), d_r(x,z), d_r(y,z)\}$$

for all x, y, and z and

$$\min\{d(c,x), d(c,y)\} + d(x,y) = \max\{d(c,x), d(c,y)\}$$

for all x and y (as well as axioms establishing the relationship between d_s for various s). The first axiom ensures $\mathbb{R}$-embeddability and the second that c is an endpoint of a model of this structure into $\mathbb{R}$, which we can take to be 0. Any such structure must have that the embedding into X is 1-Lipschitz by our requirement that $\alpha_{U_h}(x) = Lx$ whenever $h : X \to \mathbb{R}$ is an L-Lipschitz function (by [3, Thm. 1.6]), so we have that $\mathfrak{P}_r(X)$ corresponds precisely to the models of this theory.

Now we need to show that the topology on $\mathfrak{P}_r(X)$ agrees with the topology induced by taking ultraproducts. It is clear that if a net of elements of of $\mathfrak{P}_r(X)$ converges in the topology, then any ultraproduct with an ultrafilter extending the net will converge to the same structure, so we have that the ultraproduct topology is coarser than the topology on $\mathfrak{P}_r(X)$. Now consider a family $\{f_i\}_{i\in I}$ of elements of $\mathfrak{P}_r(X)$, thought of as $\mathcal{L}_X(c)$-structures, together with an ultrafilter $\mathcal{U}$ on I. Let the ultraproduct be $f_{\mathcal{U}}$. We know that this corresponds to an element of $\mathfrak{P}_r(X)$, and we will identify it with this corresponding element.

For each $s \in \operatorname{dom} f_{\mathcal{U}}$ and any finite sequence $h_1, \dots, h_n$ of continuous $\mathbb{R}$-valued functions on X, we must have that

$$f_{\mathcal{U}} \models \inf_x \max\{|d(x,c) - s|, \max_{\ell \leq n} |U_h(x) - h_\ell(f_{\mathcal{U}}(s))|\} = 0.$$

Therefore, it must be the case that for any $\varepsilon > 0$, there is a $\mathcal{U}$-large set of indices such that $\operatorname{dom} f_i$ contains an element t with distances less than ε from s such that $|h_\ell(f_i(t)) - h_\ell(f_{\mathcal{U}}(s))| < \varepsilon$ for every $i \leq \ell$. Since we can do this for any $\varepsilon > 0$ and any finite sequence of continuous $\mathbb{R}$-valued functions on X, we have by Proposition 2.2 that the family $\{f_i\}_{i\in I}$ converges along the ultrafilter $\mathcal{U}$ in the ordinary topology on $\mathfrak{P}_r(X)$.

Therefore the topologies agree, and $\mathfrak{P}_r(X)$ is compact. □

3. First-Order Theory of $\mathfrak{F}(X)$

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

Proposition 3.1. *For any* $r > 0$ *and* K *and* K' *in* $\mathfrak{F}(X)$ *with* $d(K,K') < r$, *the formula*

$$\delta_{K,K',r}(x) \coloneqq \min\{d_{2r}(x,K) + d_{2r}(x,K') - d_r(K,K'), r\}$$

is the distance predicate (with regards to d_s *for any* $s \geq r$*) of the set of* L *that satisfy either* $K \sqcap K' \sqsubseteq L \sqsubseteq K$ *or* $K \sqcap K' \sqsubseteq L \sqsubseteq K'$.

Proof. Let $[K,K']$ denote the set described in the proposition.

Let A be an element in the same finite distance component as K and K'. Assume without loss of generality that $A \sqcap K$ is longer than $A \sqcap K'$. Let $B = A \sqcap K$.

If $|B| < |K \sqcap K'|$, then we have that

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

and it is easy to check that $d(A, C) \geq d(A, K \sqcap K') = d(A, B) + d(B, K \sqcap K')$ for any $C \in [K, K']$. This implies that $d(A, [K, K']) = d(A, K) + d(A, K') - d(K, K')$.

If $|B| \geq |K \sqcap K'|$, then it must be the case that $K \sqsupseteq B \sqsupseteq K \sqcap K'$, which implies that $d(A, [K, K']) = d(A, B)$ and so also $d(A, [K, K']) = d(A, K) + d(A, K') - d(K, K')$.

So in either case we have that $d(A, [K, K']) = d(A, K) + d(A, K') - d(K, K')$. It is straightforward to check that the formula in the proposition is equal to $\min\{d(A, [K, K']), r\}$. □

Definition 3.2. For any K and K' in the same finite distance component, we will write $[K, K']$ for the set from Proposition 3.1, called the *interval between K and K'*.

We say that two intervals $[K, K']$ and $[L, L']$ are *isomorphic* if they correspond to the same element of $\mathfrak{P}(X)$ with K and L as basepoints.

Note that $[K, K']$ and $[L, L']$ are isomorphic if and only if they are isomorphic as $\mathcal{L}_X$-structures by an isomorphism taking K to L.

Since for each $r > 0$, $[K, K']$ is uniformly definable for K and K' with $d(K, K') \leq r$, the first-order theory of the structure $\mathfrak{F}(X)$ ensures that similar sets exist in any $M \equiv \mathfrak{F}(X)$.

Corollary 3.3. *For any $M \equiv \mathfrak{F}(X)$ and any K and K' in M with $d(K, K') < r$, the formula $\delta_{K,K',r}(x)$ is the distance predicate (with regards to the metric d_s for any $s \geq r$) of a set that is isometric to a closed subset of $[0, d(K, K')]$ with K and K' as endpoints.*

Proof. This is a first-order property; specifically, $[K, K']$ with the constant c assigned to K is a model of T_r, the theory of elements of $\mathfrak{P}_r(X)$. □

Proposition 3.4. *If M is any model of $T(X)$, then for any $K, K', L \in M$ with pairwise finite distance, there is a unique point $A \in [K, K']$ such that $d(L, A) = d(L, [K, K'])$.*

Proof. We will show that in $\mathfrak{F}(X)$, for any K, K', L with pairwise distance $< r$, the formula $d(x, L) - d(L, [K, K'])$ is, inside $[K, K']$, the distance predicate of a singleton. This property is preserved under ultraproducts for every $r > 0$, so the required statement will follow.

Fix $K, K', L \in \mathfrak{F}(X)$ with pairwise finite distance. There are two cases:

Case 1. Either $K \sqcap L \in [K, K']$ or $K' \sqcap L \in [K, K']$. Note that either possibility implies that $L \sqsupseteq K \sqcap K'$, so either possibility implies the other. Assume without loss of generality that $|K \sqcap L| \geq |K' \sqcap L|$. This implies that $K' \sqcap L = K \sqcap K'$. We have immediately that $d(L, [K, K']) \leq d(L, K \sqcap L)$. For any element A of $[K, K \sqcap K']$, we clearly have that $d(L, A) = d(L, K \sqcap L) + d(K \sqcap L, A)$, so $d(K \sqcap L, A) = d(L, A) - D(L, K \sqcap L)$. For any element B of $[K', K \sqcap K']$, we have that $d(L, B) = d(L, K \sqcap L) + d(K \sqcap L, K \sqcap K') + d(K \sqcap K', B)$, so $d(B, K \sqcap L) = d(B, K \sqcap K') + d(K \sqcap K', K \sqcap L) = d(L, B) - d(L, K \sqcap L)$. In either case, we have that for any $C \in [K, K']$, $d(C, K \sqcap L) = d(L, C) - d(L, K \sqcap L)$, as required.

Case 2. Neither $K \sqcap L \notin [K, K']$ nor $K' \sqcap L \in [K, K']$. This implies that $K \sqcap L \sqsubset K \sqcap K'$, which means that $K \sqcap L = K' \sqcap L$. Therefore, for any $A \in [K, K']$, we have that $d(L,A) = d(L,K \sqcap K') + d(K \sqcap K',A)$, whence $d(A,K \sqcap K') = d(L,A) - d(L,K \sqcap K')$, as required. □

Note, though, that the map taking L to the nearest point A in $[K,K']$ cannot be a definable function, since it is only well defined inside the finite distance component of $[K,K']$, which is a co-type-definable set. It is, however, representable as a family of partial definable functions with domains containing $[K,K']^{<r}$ for each $r > 0$.

Definition 3.5. In any model M of $T(X)$, a *finite tree* is a set which can be written as a union of a finite sequence $\{[K_i,L_i]\}_{i<n}$ with the property that for each $i < n$ with $i > 0$, $[K_i,L_i]$ is not disjoint from $\bigcup_{j<i}[K_j,L_j]$.

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.

Lemma 3.6. *If R is a finite tree, then for any $K,L \in R$, $[K,L] \subseteq R$.*

Proof. We will proceed by induction. First assume that R is $[A,B]$. In $\mathfrak{F}(X)$, for any given $r > 0$, we have for all A,B with $d(A,B) < r$, that any $K,L \in [A,B]$ has $[A,B] \subseteq [K,L]$. Since $[K,L]$ and $[A,B]$ are uniformly definable (as long as $d(K,L) < r$ and $d(A,B) < r$), this is a first-order fact and will still be true in any model of $T(X)$.

Assume that the assertion is true for all finite trees that are unions of $n \geq 1$ intervals. Let R be a finite tree that is the union of n intervals, and let $[A,B]$ be some interval with $R \cap [A,B]$ non-empty. If $K,L \in R$ or $K,L \in [A,B]$, then we have the assertion immediately, so assume without loss of generality that $K \in R$ and $L \in [A,B]$. Let C be some element of $R \cap [A,B]$. By the induction hypothesis, we have that $[K,C] \subseteq R$ and $[C,L] \subseteq [A,B]$.

In $\mathfrak{F}(X)$, it is easy to check that for any three elements E, F, and G in the same finite distance component, $[E,G] \subseteq [E,F] \cup [F,G]$. For any $r > 0$, this is a first-order fact for E, F, and G with pairwise distance less than r, so it will hold in any model of $T(X)$ as well.

This implies that $[K,L] \subseteq [K,C] \cup [C,L] \subseteq R \cup [A,B]$.

Therefore, by induction, we have that the result holds for all finite trees. □

Corollary 3.7. *For any $M \models T(X)$ and any finite n-tuple $\bar{a} \in M$ with pairwise finite distances, $R = \{a_0\} \cup [a_0,a_1] \cup \cdots \cup [a_{n-2},a_{n-1}]$ is the intersection of all finite trees containing $\bar{a}$.*

Proof. Clearly R is a finite tree, so the intersection of all finite trees containing $\bar{a}$ must be a superset of R. By Lemma 3.6, R must be a subset of any finite tree containing $\bar{a}$, so we are done. □

Definition 3.8. For any model $M \models T(X)$ and any 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 $\bar{a}$.

Proposition 3.9. *For any model M of $T(\mathfrak{F}(X))$, any finite tree R in M, and any $K \in M$ in the same finite distance component of M, there is a unique element $L \in R$ with $d(K,L) = d(K,R)$.*

Proof. Proceed by induction on the number of intervals in the finite tree. If the finite tree is a single interval, then by Proposition 3.4 the result follows.

Suppose that we know the result for all finite trees that are unions of n intervals. Let R be a finite tree that is the union of n intervals and let $[A, B]$ be some interval with $R \cap [A, B]$ non-empty. Clearly $d(K, R \cup [A, B]) = \min\{d(K, R), d(K, [A, B])\}$. By the induction hypothesis, there is a unique point $L_0 \in R$ with $d(K, L_0) = d(K, R)$ and a unique point $L_1 \in [A, B]$ with $d(K, L_1) = d(K, [A, B])$. If $d(K, L_0) \neq d(K, L_1)$, we are done, so assume that $d(K, L_0) = d(K, L_1)$. Again by the induction hypothesis, there is some $L_2 \in [L_0, L_1]$ such that $d(K, L_2) = d(K, [L_0, L_1])$. If $L_0 \neq L_1$, then we must have, by the uniqueness of L_2, that $d(K, L_2) < d(K, L_0) = d(K, L_1)$, but by Lemma 3.6, we have that $L_2 \in R \cup [A, B]$. This implies that either $d(K, R) < d(K, L_0)$ or $d(K, [A, B]) < d(K, L_1)$, which is absurd. Therefore there must be some unique L in $R \cup [A, B]$ such that $d(K, L) = d(K, R \cup [A, B])$.

Therefore, by induction, we have that the assertion is true for all finite trees. □

Proposition 3.10. *For any ultrafilter $\mathcal{U}$ on index set I, if $\{a_i\}_{i\in I}$ and $\{b_i\}_{i\in I}$ are families of elements of $\mathfrak{F}(X)$ whose corresponding elements a and b in $\mathfrak{F}(X)^{\mathcal{U}}$ satisfy $d(a, b) < \infty$, then $[a, b]$ in the limit corresponds to the element of $\mathfrak{P}(X)$ that is the limit along $\mathcal{U}$ of the elements in $\mathfrak{P}(X)$ corresponding to the $[a_i, b_i]$'s.*

Proof. This is immediate from Propositions 2.3 and 3.1 (as well as the fact that ultraproducts commute with restricting to definable sets). □

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

Lemma 3.11 (Parallel Paths)**.** *(∂ open.) For any $f \in \mathfrak{P}(X)$, 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 a path $g \in \mathfrak{P}(X)$ such that $g(0) = x$ and $(f, g) \in U_{V,\varepsilon}$.*

Proof. For any entourage $W \subseteq X^2$ and any $z \in X$, let $W(z) = \{w : (w, z) \in W\}$. Obviously $W(z)$ is an open neighborhood of z.

Find an 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.)

Let F be a finite subset of $\mathrm{dom}\, f$ such that $0 \in F$ and $d_H(F, \mathrm{dom}\, f) < \frac{1}{2}\delta$. (Such a set always exists by compactness.) Let $\{r_i\}_{i\leq n}$ be an increasing enumeration of F (with $r_0 = 0$). Fix $\gamma > 1$ small enough that $(\gamma - 1)|f| < \frac{1}{2}\delta$. For each $i < n$, let $d(i) = d(f(r_i), f(r_{i+1}))$.

Let $A_i = W(f(r_i))$ for each $i \leq n$. For each $i < n$,

- if $f(r_i) = f(r_{i+1})$, let B_i and C_{i+1} both equal X and
- if $f(r_i) \neq f(r_{i+1})$, find open neighborhoods $B_i \ni f(r_i)$ and $C_{i+1} \ni f(r_{i+1})$ such that $\mathrm{cl}(B_i)^{\leq\gamma^{-1}d(i)} \cap \mathrm{cl}(C_{i+1}) = \varnothing$. (This is always possible in any compact topometric space.)

Let $D_0 = A_0 \cap B_0$, $D_i = A_i \cap B_i \cap C_i$ for each $0 < i < n$, and $D_n = A_n \cap C_n$. Note that each D_i is an open neighborhood of $f(r_i)$. What we have guaranteed at this point is that for each $i < n$, if $y \in D_i$ and $z \in D_{i+1}$, then $\partial(y, z) \geq \gamma^{-1}d(i)$.

Let $E_n = D_n$. For each $i < n$, let $E_i = D_i \cap E_{i+1}^{<\gamma d(i)}$. Note that each E_i is an open neighborhood of $f(r_i)$ (since d is an open metric). Now note that for each $i < n$, for any $y \in E_i$, there is a $z \in E_{i+1}$ such that $\partial(y, z) < \gamma d(i)$. Since y and z are also in D_i and D_{i+1}, respectively, we also have that $\partial(y, z) \geq \gamma^{-1}d(i)$.

So let $O = E_0$. For any $x \in O$, by construction, we can find a sequence $\{x_i\}_{i\leq n}$ such that

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

We are note 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 $\mathfrak{P}(X)$ with domain γF with the property that for each $i \leq n$, $g(\gamma r_i) = x_i$. We want to show that $g \in U_{f,V,\varepsilon}$.

For each γr_i in $\operatorname{dom} g$, 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 $s \in \operatorname{dom} f$, 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 $f(s) \in (\operatorname{cl} W(g(\gamma r_i)))^{\leq\delta} \subseteq V(g(\gamma r_i))$, as required.

Therefore $g \in U_{f,V,\varepsilon}$. Finally, $|g| \leq \gamma|f| < |f| + \varepsilon$. □

Definition 3.12. For any model $M \models T(X)$ and any $a \in M$, we write $\operatorname{tp}_X(a)$ for the unique $x \in X$ with the property that for every continuous $f : X \to \mathbb{R}$, $f(x) = U_f^M(a)$.

Note that $\operatorname{tp}_X(a)$ 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 3.13. *(∂ open.) For any model $M \models T(X)$, any $b \in M$, any $f \in \mathfrak{P}(X)$ with $f(0) = 0$, 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 such that for any $i < j < \kappa$, $[b, c_i] \cap [b, c_j] = \{b\}$.*

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

Lemma 3.11 implies that for each entourage V and $\varepsilon > 0$, we can find an open neighborhood $O \ni f(0)$ such that for any $x \in O$, there is a path $g \in \mathfrak{P}(X)$ such that $g(0) = x$ and $(f, g) \in U_{V,\varepsilon}$. This implies that for a $\mathcal{F}$-large set of i, we can find $L_i^n \in \mathfrak{F}(X)$ for each $n < \omega$ such that

- $[K_i, L_i^n]$ corresponds to some g in $\mathfrak{P}(X)$ with $g(0) = (K_i)_X(0)$ and $(f, g) \in U_{V,\varepsilon}$ for every $n < \omega$ and
- for any $n < k < \omega$, $[K_i, L_i^n] \cap [K_i, L_i^k] = \{K_i\}$.

Since we can do this for any entourage V and $\varepsilon > 0$, we have, by compactness, that there is an elementary extension $C \succ \mathfrak{F}(X)$ and a family $\{B^n\}_{n<\omega}$ of elements of C such that

- $[a, B^n]$ corresponds to $f \in \mathfrak{P}(X)$ for every $n < \omega$ and
- for any $n < k < \omega$, $[a, B^n] \cap [a, B^k] = \{K_i\}$.

Since $b \equiv a$, the required elementary extension of N must exist as well. □

Lemma 3.14. *(∂ open.) Let $\mathcal{U}$ be the monster model of $T(X)$. 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 tuples of elements of $\mathcal{U}$ 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 $f : A \cong B$ such that for each i, j, $f(a_i^j) = b_i^j$. Then for any $c \in \mathcal{U}$, there exists an $e \in \mathcal{U}$ 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 $g : Ac \to Be$ 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 $g : A \cup \mathrm{ccl}(\bar{a}^i c) \cong B \cup \mathrm{ccl}(\bar{b}^i e)$ 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 $e \in \mathcal{U}$ 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 3.9, 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 $\mathfrak{P}(X)$ corresponding to $[c', c]$. By assumption, we have that $e' := f(c') \in B$ has $\mathrm{tp}_X(e') = \mathrm{tp}_X(c')$, so by Lemma 3.13, there is a family $\{e_i\}_{i<(2^{\aleph_0})^+}$ of elements of $\mathcal{U}$ such that for each $i < (2^{\aleph_0})^+$, $[e', e_i]$ corresponds to h in $\mathfrak{P}(X)$ and for each $i < j < (2^{\aleph_0})^+$, $[e', e_i] \cap [e', e_j] = \{e'\}$. Since the cardinality of $\mathrm{ccl}(\bar{b}^i)$ is at most $2^{\aleph_0}$, by the pigeonhole principle, there must be some $i < (2^{\aleph_0})^+$ 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 $\mathfrak{P}(X)$, we have that g is an $\mathcal{L}_X$-isomorphism. Finally, we clearly have that $g(c) = e$. ☐

Proposition 3.15. *(∂ open.) For any finite tuple $\bar{a}$ in any model $M \models T(X)$, $\mathrm{tp}(\bar{a})$ is uniquely determined by the partitioning of $\bar{a}$ into finite distances 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 3.14 and a back-and-forth argument. ☐

Corollary 3.16. *(∂ open.) For any $a \in M \models T(X)$, $\mathrm{tp}(a)$ 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 3.15, we have that $\mathrm{tp}(a) = \mathrm{tp}(b)$. ☐

4. Stability of $T(X)$

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

Lemma 4.1. *For any model $M \models T(X)$, 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 $F := \bigcap\{[a, c]^N : c \in M, d(a, c) < \infty\}$. Since this is the intersection of a family of compact sets, it itself is compact. So in particular, 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, by compactness, there must be 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 3.9, 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 be the unique element of $[m, a]$ of minimal distance to g.

It is easy to check that in $\mathfrak{F}(X)$, for any A, B, C, the unique element of $[A, B]$ closest to C is contained in $[A, C]$, which implies that this is true for all models of $T(X)$. Therefore we have that $h \in [m, g]$. Since $[m, g]$ is in the algebraic closure of mg, we have that $[m, g] \subset M$. By Lemma 3.6, 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 4.2. *Fix $a, b, c, e \in M \models T(X)$ with pairwise finite distance. Let c' be the unique closest point on $[a, b]$ to c and e' the unique closest point on $[a, b]$ to e. If $c' \neq e'$, then $d(c, e) = d(c, c') + d(c', e') + d(e', e)$.*

Proof. It is easy to verify that this is true in $\mathfrak{F}(X)$. For any $r, \varepsilon > 0$, the statement

> for any a, b, c, e with pairwise distance $< r$, if c' is the closest point on $[a, b]$ to c and e' to e and $d(c', e') > \varepsilon$, then $d(c, e) = d(c, c') + d(c', e') + d(e', e)$

is first-order. Therefore these hold in any model of $T(X)$, which is precisely the desired conclusion. $\square$

Lemma 4.3. *For any κ, there is a model M of $T(X)$ with density character κ such that $|M| = \kappa^\omega$.*

Proof. This follows the argument in the proof of Theorem 8.10 in [5], noting that any given $\mathbb{R}$-tree embeds isometrically into a model of $T(X)$. $\square$

Proposition 4.4. *For any model $M \models T(X)$, the elements of $S_1(M)$ are precisely*

- *the realized types in M,*
- *types $p_{m,f}$ for each pair $m \in M$ and $f \in \mathfrak{P}(X)$ with $f(0) = \mathrm{tp}_X(m)$, 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 $[m, a]$ corresponds to f in $\mathfrak{P}(X)$ 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[5] *on $S_1(M)$ is given by*

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

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

- $d(q_x, q_{x'}) = \partial(x, x')$, *and*
- $d(m, q_x) = d(p_{m,f}, q_x) = \operatorname{diam} X$

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

$$|M| + \#^{\mathrm{dc}} X \leq \#^{\mathrm{dc}} S_1(M) \leq |\mathfrak{P}(X)| \cdot |M| + \#^{\mathrm{dc}} X,$$

and $T(X)$ *is strictly stable.*

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 $m \in M$. The characterizations of these types as $p_{m,f}$ and q_x for various m,f, and x follows from Lemma 4.1, Proposition 3.15, and Corollary 3.16.

For the metric 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 4.2, $d(a, b) = d(a, c) + d(c, e) + d(e, b)$, as required.

The bounds on the density character of $S_1(M)$ are obvious, so the fact that $T(X)$ is strictly stable follows from Lemma 4.3. □

5. Main Theorem

Lemma 5.1. *Any compact topometric space* (X, τ, ∂) *with open metric has finite diameter.*

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

So assume 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, there are open neighborhoods $U_1 \ni x_1$ and $V_1 \ni y_1$ such 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 $\partial(x_{i+1}, y_{i+1}) > i + 1$. 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). □

Theorem 5.2. *For any compact topometric space* (X, τ, ∂) *with an open metric, there is a stable continuous first-order theory* T *such that* $S_1(T)$ *is isomorphic to* (X, τ, ∂).

Proof. If X has a single point, then the theory of a one-point structure suffices, so assume that X has more than one point. By Lemma 5.1, X has finite diameter, so we can form the theory $T(X)$. There is clearly a continuous 1-Lipschitz map f from $S_1(T)$ to X. We have by Corollary 3.16 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 $\mathfrak{F}(X)$ with $\operatorname{tp}_X(K) = a$ and $\operatorname{tp}_X(L) = b$ such that $d_{\operatorname{diam} X}(K, L) = d(K, L) = \partial(a, b)$. Therefore f is an isometry as well. Finally, by Proposition 4.4, $T(X)$ 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 T superstable?

Clearly if X is not CB-analyzable,[6] then any such T cannot be ω-stable or totally transcendental, but it is also possible that this is the only obstruction.

Question 5.3. *If (X, τ, ∂) is a compact topometric space with an open 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] Itaï Ben Yaacov. On perturbations of continuous structures. *Journal of Mathematical Logic*, 08(02):225–249, 2008.

[2] Itaï Ben Yaacov. Topometric spaces and perturbations of metric structures. *Logic and Analysis*, 1(3):235, July 2008.

[3] Itaï Ben Yaacov. Lipschitz functions on topometric spaces. *Journal of Logic and Analysis*, 10 2010.

[4] 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.

[5] Sylvia Carlisle and C. Ward Henson. Model theory of $\mathbb{R}$-trees. *Journal of Logic and Analysis*, 12, December 2020.

[6] James Hanson. Strongly Minimal Sets and Categoricity in Continuous Logic. *arXiv e-prints*, page arXiv:2011.00610, November 2020.

Email address: `jehanson2@wisc.edu`

Department of Mathematics, University of Wisconsin–Madison, 480 Lincoln Dr., Madison, WI 53706

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