Model theory of $\mathbb{R}$-trees

Sylvia Carlisle

C. Ward Henson

Abstract: We show the theory of pointed $\mathbb{R}$-trees with radius at most r is axiomatizable in a suitable continuous signature. We identify the model companion $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ of this theory and study its properties. In particular, the model companion is complete and has quantifier elimination; it is stable but not superstable. We identify its independence relation and find built-in canonical bases for non-algebraic types. Among the models of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ are $\mathbb{R}$-trees that arise naturally in geometric group theory. In every infinite cardinal, we construct the maximum possible number of pairwise non-isomorphic models of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$; indeed, the models we construct are pairwise non-homeomorphic. We give detailed information about the type spaces of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$. Among other things, we show that the space of 2-types over the empty set is nonseparable. Also, we characterize the principal types of finite tuples (over the empty set) and use this information to conclude that $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ has no atomic model.

1 Introduction

Continuous logic is an extension of classical first order logic used to study the model theory of structures based on metric spaces. In this paper, we use continuous logic as presented in Ben Yaacov, Berenstein, Henson and Usvyatsov [4] and Ben Yaacov and Usvyatsov [6] to study the model theory of $\mathbb{R}$-trees.

An $\mathbb{R}$-tree is a metric space T such that for any two points $a, b \in T$ there is a unique arc in T from a to b, and that arc is a geodesic segment (ie, an isometric copy of some closed interval in $\mathbb{R}$). These spaces arise naturally in geometric group theory, for example: the asymptotic cone of a hyperbolic finitely generated group is an $\mathbb{R}$-tree.

An $\mathbb{R}$-tree may be unbounded, while the existing full treatments of continuous model theory are restricted to bounded structures. With this in mind, we consider pointed trees, choose a real number $r > 0$, and axiomatize the theory $\mathbb{R}T_r$ of pointed $\mathbb{R}$-trees of radius at most r in a suitable continuous signature.

We then define the notion of *richly branching* and axiomatize the theory $rb\mathbb{R}T_r$ of the class of richly branching pointed $\mathbb{R}$-trees with radius r. We prove that the models of $rb\mathbb{R}T_r$ are exactly the existentially closed models of $\mathbb{R}T_r$; thus $rb\mathbb{R}T_r$ is the model companion of $\mathbb{R}T_r$. Next, we investigate some model theoretic properties of $rb\mathbb{R}T_r$, showing that it is complete and has quantifier elimination. In particular, that means $rb\mathbb{R}T_r$ is the model completion of $\mathbb{R}T_r$. Further, we prove that $rb\mathbb{R}T_r$ is stable but not superstable and identify its model-theoretic independence relation. We characterize the principal types of $rb\mathbb{R}T_r$, and show that this theory has no atomic model. Finally, we show that $rb\mathbb{R}T_r$ is highly non-categorical. In fact, for any density character this theory has the maximum possible number of pairwise non-isomorphic models; indeed, the models we construct are pairwise non-homeomorphic. We also give examples of richly branching $\mathbb{R}$-trees which come from the literature, including some that will be familiar to geometric group theorists.

In the remainder of this introduction we detail the contents of each section of this paper.

In Sections 2 and 3 we provide background concerning $\mathbb{R}$-trees and continuous logic, respectively. In Section 4 we specify a continuous signature L suitable for the class of pointed $\mathbb{R}$-trees of radius at most r, and axiomatize this class of L-structures; the theory of the class is denoted $\mathbb{R}T_r$.

In Section 5 we discuss definability of certain sets and functions in $\mathbb{R}T_r$. In Section 6 we show $\mathbb{R}T_r$ has amalgamation over substructures. This plays an important role in many of the primary results in this paper.

In Section 7 we introduce the class of richly branching pointed $\mathbb{R}$-trees with radius r and axiomatize this class. The associated theory is denoted $rb\mathbb{R}T_r$. We then show that $rb\mathbb{R}T_r$ is the model companion of $\mathbb{R}T_r$. The theory $rb\mathbb{R}T_r$ is the main object of study in this paper.

In Sections 8 and 9 we verify the main model-theoretic properties of $rb\mathbb{R}T_r$. We show that this theory is complete and admits quantifier elimination. We characterize its types over sets of parameters and use this to show $rb\mathbb{R}T_r$ is κ-stable if and only if $\kappa = \kappa^\omega$; hence this theory is strictly stable (stable but not superstable). We also show that $rb\mathbb{R}T_r$ is not a small theory; indeed, its space of 2-types over $\emptyset$ has metric density character 2^ω. (The space of 1-types over $\emptyset$ is isometric to the real interval $[0, r]$

with the usual absolute value metric.) We give a simple geometric characterization of the independence relation of rb$\mathbb{R}$T$_r$. Finally, we show that non-algebraic types have built-in canonical bases (ie, these bases are sets of ordinary elements in models of rb$\mathbb{R}$T$_r$ and do not require the introduction of imaginaries).

In Section 10 we discuss some models of rb$\mathbb{R}$T$_r$ that have been constructed within the theory of $\mathbb{R}$-trees by Dyubina and Polterovich [12, 13] and some other models that arise in geometric group theory.

In Section 11 we use amalgamation constructions to build large families of models of rb$\mathbb{R}$T$_r$ and to characterize its isolated types over $\emptyset$. For each infinite cardinal κ, we show there are 2^{κ}-many pairwise non-isomorphic models of rb$\mathbb{R}$T$_r$ of density character κ. This is the maximum possible number of models, and the models we construct are, in fact, pairwise non-homeomorphic. We show that rb$\mathbb{R}$T$_r$ has very few isolated n-types over $\emptyset$ and conclude that it has no atomic model (equivalently, it has no prime model).

In Section 12 we briefly discuss how the results in this paper could be obtained for the full class of pointed $\mathbb{R}$-trees (ie, without imposing a boundedness requirement).

2 $\mathbb{R}$-trees

In this section we give some background concerning $\mathbb{R}$-trees.

2.1 Definition   An *arc* in a metric space M is the image of a homeomorphism γ from an interval $[0, r] \subseteq \mathbb{R}$ into M for some $r \geq 0$. A *geodesic segment* in a metric space M is the image of an isometric embedding $\gamma\colon [0, r] \to M$. We say that such an arc or geodesic segment is *from $\gamma(0)$ to $\gamma(r)$*. A metric space M is called *geodesic* if for every $a, b \in M$ there is at least one geodesic segment in M from a to b.

2.2 Fact   A complete metric space M is a geodesic space if and only if for any $x, y \in M$ there exists a midpoint z between x and y. That is, there exists z such that:

$$d(x, z) = \frac{d(x, y)}{2} \quad \text{and} \quad d(y, z) = \frac{d(x, y)}{2}$$

Proof   See Bridson and Haefliger [7, I.1.4, page 4].   □

2.3 Definition An $\mathbb{R}$-*tree* is a metric space M such that for any two points $a, b \in M$ there is a unique arc from a to b, and that arc is a geodesic segment.

In an $\mathbb{R}$-tree, $[a, b]$ denotes the unique geodesic segment from a to b. Since metric structures are required to be based on complete metric spaces, it is a helpful fact that the completion of an $\mathbb{R}$-tree is an $\mathbb{R}$-tree (see Chiswell [8, Lemma 2.4.14]).

Let M be an $\mathbb{R}$-tree and $a \in M$. Call the connected components of $M \setminus \{a\}$ *branches* at a. Let the *degree* of a point $a \in M$ be the cardinal number of branches at a. If there are three or more branches at $a \in M$, then we call a a *branch point*. The *height* of a branch β at a is $\sup\{d(a, x) | x \in \beta\}$ if that supremum exists, and is ∞ otherwise. A *subtree* of M is any subspace of M that is itself an $\mathbb{R}$-tree. A *ray* in an $\mathbb{R}$-tree is an isometric copy of $\mathbb{R}^{\geq 0}$. If $a \in M$, a *ray at a* is a ray so that the image of 0 under the isometric embedding of $\mathbb{R}^{\geq 0}$ into M is a.

The following lemmas and definitions collect some straightforward facts about $\mathbb{R}$-trees used in this paper. For helpful pictures and more facts about $\mathbb{R}$-trees see [8].

2.4 Lemma *If M is an $\mathbb{R}$-tree and $a, b, c \in M$, then:*
  (1) $d(a, b) + d(b, c) = d(a, c) + 2\operatorname{dist}(b, [a, c])$.
  (2) $b \in [a, c]$ *if and only if* $d(a, c) = d(a, b) + d(b, c)$.
  (3)   *For b distinct from a and c, we have $b \in [a, c]$ if and only if a and c are in different branches at b.*

Proof Statement (1) follows from [8, Lemma 2.1.2], (2) is proved in [8, Lemma 1.2.2] and (3) comes from [8, Lemma 2.2.2].   □

2.5 Lemma *If M is an $\mathbb{R}$-tree and E_1, E_2 are disjoint, closed, non-empty subtrees of M, then there exists a unique shortest geodesic segment $[u, v]$ with $u \in E_1$ and $v \in E_2$. Moreover, for all $b \in E_1$ and $c \in E_2$, the geodesic segment from b to c must contain $[u, v]$.*

Proof This is [8, Lemma 2.1.9].   □

The preceding lemma directly implies the following fact, used often in this paper: Given an $\mathbb{R}$-tree M, a closed subtree E and a point $a \in M$, there is a unique point e in E closest to a. In that situation, one has that $\operatorname{dist}(a, E) = d(a, e)$ and also that e is on the segment $[a, b]$ for every point $b \in E$. Equivalently, given any $b \in E$, if $\gamma \colon [0, d(a, b)] \to M$ is an isometry and has the geodesic segment $[a, b]$ as its image, then $e = \gamma(s)$, where $s = d(a, e)$ is the least value of $t \in [0, d(a, b)]$ for which $\gamma(t) \in E$.

2.6 Definition Let $x_0, x_1, x_2, \ldots, x_n$ be points in an $\mathbb{R}$-tree M and $\gamma \colon [0, d(x_0, x_n)] \to M$ the isometric embedding with $\gamma(0) = x_0$ and $\gamma(d(x_0, x_n)) = x_n$ that has image equal to the geodesic segment $[x_0, x_n]$. If for each $i = 0, \ldots, n$ we have $x_i = \gamma(a_i)$ where $0 = a_0 \leq a_1 \leq \cdots \leq a_n = d(x_0, x_n)$, then we write $[x_0, x_n] = [x_0, x_1, \ldots, x_n]$, and call $[x_0, x_1, \ldots, x_n]$ a *piecewise segment*.

In other words, if $x_0, x_1, \ldots, x_n$ are elements of $[x_0, x_n]$ listed in increasing order of distance from x_0, then we write $[x_0, x_n] = [x_0, x_1, \ldots, x_n]$ for this *piecewise segment*. Note that we also know $[x_1, x_n] = \bigcup_{i=1}^{n-1} [x_i, x_{i+1}]$, and by Lemma 2.1.4 in Chiswell [8] we have that $[x_1, x_n] = [x_1, x_2, \ldots, x_n]$ if and only if $d(x_1, x_n) = \sum_{i=1}^{n-1} d(x_i, x_{i+1})$.

2.7 Lemma *Let E be a closed subtree of an $\mathbb{R}$-tree M and let $a, b \in M$. Let $e_a \in E$ be the unique closest point to a, and let $e_b \in E$ be the unique closest point to b. If $e_a \neq e_b$, then*

$$d(a, b) = d(a, e_a) + d(e_a, e_b) + d(b, e_b).$$

That is, $[a, e_a, e_b, b]$ is a piecewise segment.

Proof Follows from Lemmas 2.4 and 2.5. $\qquad\qquad\qquad\qquad\qquad\qquad\square$

Recall that (X, d) is a *pseudometric space* if $d \colon X^2 \to \mathbb{R}^{\geq 0}$ is symmetric, satisfies the triangle inequality, and has $d(x, x) = 0$ for all $x \in X$. Its quotient metric space is obtained by identifying $x, y \in X$ iff $d(x, y) = 0$.

2.8 Definition (*Gromov product*) For a pseudometric space (X, d) and $x, y, w \in X$, define:

$$(x \cdot y)_w = \frac{1}{2}[d(x, w) + d(y, w) - d(x, y)]$$

It follows easily from Lemmas 2.4 and 2.5 that in an $\mathbb{R}$-tree, the Gromov product $(x \cdot y)_w$ computes the distance from w to the geodesic segment $[x, y]$.

2.9 Definition Let $\delta \geq 0$. A pseudometric space (X, d) is δ-*hyperbolic* if, for all $x, y, z, w \in X$:

$$\min\{(x \cdot z)_w, (y \cdot z)_w\} - \delta \leq (x \cdot y)_w$$

Note that if (X, d) is a pseudometric space, then the quotient metric space of (X, d) is δ-hyperbolic if and only if (X, d) is δ-hyperbolic.

In this paper we are nearly always interested in 0-hyperbolic spaces, and it is sometimes useful to have the following alternate characterization:

2.10 Lemma *A pseudometric space (X, d) is 0-hyperbolic if and only if it satisfies the following 4-point condition: for all $x, y, z, t \in X$,*

$$d(x, y) + d(z, t) \leq \max\{d(x, z) + d(y, t), d(y, z) + d(x, t)\}.$$

Proof See Chiswell [8, Lemma 1.2.6] with $\delta = 0$. □

2.11 Lemma *If M is a geodesic metric space, then M is 0-hyperbolic if and only if, given any $a, b, c \in M$ and geodesic segments $[a, b], [b, c]$, and $[c, a]$, the segment $[a, b]$ is contained in $[b, c] \cup [c, a]$.*

Proof See the proof of Proposition III.H.1.22 in Bridson and Haefliger [7]. □

2.12 Definition If $A \subseteq M$ is a subset of the $\mathbb{R}$-tree M, let E_A denote the smallest subtree containing A. We call this the $\mathbb{R}$-*tree spanned by A*. Note that

$$E_A = \bigcup\{[a_1, a_2] \mid a_1, a_2 \in A\}.$$

The closure $\overline{E_A}$ of E_A is the smallest closed subtree containing A.

Note that if A is finite, then $E_A = \overline{E_A}$ and E_A is complete and has finite diameter.

2.13 Definition An $\mathbb{R}$-tree M is *finitely spanned* if there exists a finite subset $A \subseteq M$ such that $M = E_A$.

2.14 Lemma (1) *A metric space is an $\mathbb{R}$-tree if and only if it is 0-hyperbolic and geodesic.*
  (2) *Any 0-hyperbolic metric space embeds isometrically in an $\mathbb{R}$-tree.*
  (3) *Let (X, d) be a 0-hyperbolic metric space. For $i = 1, 2$, suppose $f_i \colon (X, d) \to (M_i, d_i)$ are isometric embeddings into $\mathbb{R}$-trees, and let E_i be the smallest subtree of M_i containing $f_i(X)$. Then there is a unique isometry g from E_1 onto E_2 such that $g \circ f_1 = f_2$.*

Proof (1), (2) These are from Roe [15, Proposition 6.12].

(3) Let (M, d) be any $\mathbb{R}$-tree extending (X, d) and let E be the smallest subtree of M containing X.

For any points $u, u' \in E$ there exist $x, y, x', y' \in X$ such that $u \in [x, y], u' \in [x', y']$. (See Definition 2.12.) The main point is to show that $d(u, u')$ is determined by

the distances $d(x, u), d(x', u')$ together with the pairwise distances between points in $\{x, y, x', y'\}$.

Fix u, u', x, y, x', y' as above. By Lemma 2.5 there exist $Y \in [x, y]$ and $Y' \in [x', y']$ such that for any $v \in [x, y], v' \in [x', y']$ we have:

(D) $$d(v, v') = d(v, Y) + d(Y, Y') + d(Y', v')$$

(In particular, $[Y, Y']$ is the shortest geodesic segment among those between a point in $[x, y]$ and a point in $[x', y']$.)

Applying (D) to x, x' and to y, y' yields

$$d(x, x') = d(x, Y) + d(Y, Y') + d(Y', x')$$
and $$d(y, y') = d(y, Y) + d(Y, Y') + d(Y', y'),$$

from which follows

$$2d(Y, Y') = d(x, x') + d(y, y') - d(x, y) - d(x', y')$$

which shows that $d(Y, Y')$ is determined by the values of d on x, y, x', y'.

Similarly, applying (D) to x, y' and to x, x' yields that $d(x, Y)$ is determined by the values of d on x, y, x', y'. Likewise, applying (D) to x', y and to x', x yields the same conclusion for $d(x', Y')$.

Finally, applying (D) to the original pair u, u' yields

$$\begin{aligned} d(u, u') &= d(u, Y) + d(Y, Y') + d(Y', u') \\ &= |d(x, Y) - d(x, u)| + d(Y, Y') + |d(x', Y') - d(x', u')| \end{aligned}$$

which gives the desired conclusion.

To construct the needed isometry g, there is an obvious definition from a segment of the form $[f_1(x), f_1(y)]$ in M_1 into M_2, for each $x, y \in X$, by taking g to be the unique isometry from $[f_1(x), f_1(y)]$ in M_1 onto $[f_2(x), f_2(y)]$ in M_2 that takes $f_1(x)$ to $f_2(x)$ and $f_1(y)$ to $f_2(y)$. What is proved above shows that the union of all such maps is a well defined isometry from E_1 onto E_2 that satisfies $g \circ f_1 = f_2$. Every isometry from E_1 to E_2 that satisfies $g \circ f_1 = f_2$ must agree with this map g. $\square$

Given a, b, c in an $\mathbb{R}$-tree, there is a unique point Y so that $[a, b] \cap [a, c] = [a, Y]$. In Chiswell [8] after Lemma 2.1.2, it is shown that this Y is also the unique point so that $[b, a] \cap [b, c] = [b, Y]$ and $[c, a] \cap [c, b] = [c, Y]$, and that in fact $\{Y\} = [a, b] \cap [b, c] \cap [a, c]$.

2.15 Notation If a, b, c are points in an $\mathbb{R}$-tree, we denote by $Y(a, b, c)$ the unique point such that $\{Y\} = [a, b] \cap [b, c] \cap [a, c]$. This point will be denoted simply by Y when a, b, c are understood.

That an $\mathbb{R}$-tree is 0-hyperbolic tells us that for any 3 points a, b, c the segment $[a, b]$ is contained in $[b, c] \cup [c, a]$, by 2.11 and 2.14. Thus, the subtree E spanned by a, b, c is comprised of the segments $[a, Y]$, $[b, Y]$ and $[c, Y]$, which share only the point $Y = Y(a, b, c)$. Therefore we have two possibilities: either Y is one of a, b, c and E is the segment connecting the other two; or $Y \notin \{a, b, c\}$ and E is "Y-shaped", explaining the notation introduced in 2.15 above. (See Chiswell [8, 2.1.2].)

2.16 Definition Let M be an $\mathbb{R}$-tree. If $c \in M$ is such that there do *not* exist $a, b \in M \setminus \{c\}$ with $c \in [a, b]$, then c is called an *endpoint* of M. Equivalently, an endpoint is a point with degree at most one.

2.17 Lemma *If an $\mathbb{R}$-tree M is finitely spanned and C is the set of endpoints of M, then:*
  (1) *If B spans M then $C \subseteq B$.*
  (2) *The set C spans M.*
Thus, C is finite, and it is the unique smallest set that spans M.

Proof Let M be a finitely spanned $\mathbb{R}$-tree. Let D be the diameter of M. Let B be a set that spans M.

Proof of (1): Assume there is an endpoint $c \in M$ not contained in B. Then there must exist $a, b \in B$ such that $c \in [a, b]$ and $c \neq a, b$. But, this is a contradiction because c is an endpoint.

Proof of (2): Let $a \in M$. Let S_a be the set of all segments $[b, c] \subseteq M$ such that $a \in [b, c]$ and order S_a by inclusion. This is a partial ordering on S_a. Let $\{[b_i, c_i] \mid i \in \alpha\}$ be a chain in this partial ordering, where α is a cardinal. Let I be the closure of $\bigcup_{i \in \alpha}[b_i, c_i]$. Then I is a geodesic segment in M. Clearly $a \in I$, and the length of I is at most D. Therefore $I \in S_a$, and I is an upper bound for the chain. The chain was arbitrary, so any chain has an upper bound. Therefore, by Zorn's Lemma there exists a maximal element of S_a. Let $[b_a, c_a]$ denote such a maximal element. The elements b_a and c_a must be endpoints of M. Say, for instance, that b_a is not an endpoint. Then there exist $e, f \in M$ such that $b_a \in [e, f]$ and $b_a \neq e, f$, and then either $[e, c_a]$ or $[f, c_a]$ will properly contain $[b_a, c_a]$. This would mean $[b_a, c_a]$ was not maximal in S_a. Therefore, for each $a \in M$, there exist endpoints b_a and c_a so that $a \in [b_a, c_a]$. So, M is spanned by the set of its endpoints, and this spanning set is as small as possible by (1). $\square$

As noted above, if (X, d) is a $\mathbb{R}$-tree, then so is its completion $(\overline{X}, d)$. The next result gives information about the structure of $\overline{X}$.

2.18 Lemma *Let (X, d, p) be a pointed $\mathbb{R}$-tree and $(\overline{X}, d, p)$ its completion, and suppose $x \in \overline{X} \setminus X$. Then:*

(1) *there exist (x_n) in X converging to x such that $[p, x_1, \ldots, x_n, x]$ is a piecewise segment in $\overline{X}$ for all $n \geq 1$; and*

(2) *x is an endpoint in $\overline{X}$.*

Proof (1) Given $x \in \overline{X} \setminus X$, let (y_n) be a Cauchy sequence in X that converges to x in the $\mathbb{R}$-tree $\overline{X}$. Without loss of generality we may assume that $(d(y_n, x))$ is decreasing toward 0. For each n, let x_n be the closest point in $\overline{X}$ on the segment $[p, x]$ to y_n. By Lemma 2.5, x_n is on the segment $[p, y_n]$ in $\overline{X}$ and thus $x_n \in X$. Further, from the same Lemma and the assumptions on (y_n) we conclude that $[p, x_1, \ldots, x_n, x]$ is a piecewise segment in $\overline{X}$ for each n and (x_n) converges to x.

(2) If x were not an endpoint in $\overline{X}$, there would be $y \in \overline{X}$ such that p and y were in different branches in $\overline{X}$ at x. By part (1), there must exist $z \in X$ on the segment $[p, y]$ in $\overline{X}$ such that z is closer to y than x. That is, $[p, x, z, y]$ would be a piecewise segment in $\overline{X}$. But then x would be on the segment $[p, z]$, which lies entirely in X. This contradiction completes the proof. □

3 Some continuous model theory

We investigate the model theory of $\mathbb{R}$-trees using continuous logic for metric structures as presented in Ben Yaacov, Berenstein, Henson and Usvyatsov [4] and Ben Yaacov and Usvyatsov [6]. In this section we remind the reader of a few key concepts and results from those papers, and then we present a few facts about existentially closed models and model companions that are not discussed there. For the rest of this section we fix a continuous first order language L.

As explained in [4, Section 3], in continuous model theory it is required that *structures* and *models* are metrically complete. However, formulas and conditions are evaluated more generally in *pre-structures*, as explained in [4, Definition 3.3]. Further, it is shown in [4, Theorem 3.7] that the completion of a prestructure is an elementary extension. In this paper we use notation of the form $\mathcal{M} \models \theta$ only when $\mathcal{M}$ is a structure; in other words, M must be metrically complete.

Next, some reminders about saturation in continuous logic. A set $\Sigma(x_1, \ldots, x_n)$ of L-conditions (with free variables among $x_1, \ldots, x_n$) is called *satisfiable in* $\mathcal{M}$ if there exist $a_1, \ldots, a_n$ in M such that $\mathcal{M} \models E[a_1, \ldots, a_n]$ for every $E(x_1, \ldots, x_n) \in \Sigma$. Let κ be a cardinal. A model $\mathcal{M}$ of T is called κ-*saturated* if for any set of parameters $A \subseteq M$ with cardinality $< \kappa$ and any set $\Sigma(x_1, \ldots, x_n)$ of $L(A)$-conditions, if every finite subset of $\Sigma(x_1, \ldots, x_n)$ is satisfiable in $(\mathcal{M}, a)_{a \in A}$, then the entire set $\Sigma(x_1, \ldots, x_n)$ is satisfiable in $(\mathcal{M}, a)_{a \in A}$.

3.1 Proposition *For any countably incomplete ultrafilter U on I, the U-ultraproduct of a family of L-structures $(\mathcal{M}_i \mid i \in I)$ is ω_1-saturated.*

Proof See Ben Yaacov, Berenstein, Henson and Usvyatsov [4, Proposition 7.6]. □

Note that any non-principal ultrafilter on $\mathbb{N}$ is countably incomplete.

3.2 Proposition *For any cardinal κ, any L-structure $\mathcal{M}$ has a κ-saturated elementary extension.*

Proof See [4, Proposition 7.10]. □

3.3 Definition An inf-*formula* of L is a formula of the form

$$\inf_{y_1} \ldots \inf_{y_n} \varphi(x_1, \ldots, x_k, y_1, \ldots, y_n)$$

where $\varphi(x_1, \ldots, x_k, y_1, \ldots, y_n)$ is quantifier-free.

A sup-formula of L is defined similarly. These sup-formulas are the universal formulas in continuous logic. For an L-theory T, we use the notation $T_\forall$ for the set of universal sentences σ (sup-formulas with no free variables) such that the condition $\sigma = 0$ is implied by the theory T. Note that, as in classical first order logic, $T_\forall$ is the theory of the class of L-substructures of models of T.

3.4 Definition Let T be an L-theory and suppose $\mathcal{M} \models T$. We say $\mathcal{M}$ is an *existentially closed (e.c.)* model of T if, for any inf-formula $\psi(x_1, \ldots, x_m)$, any $a_1, \ldots, a_k \in M$, and any $\mathcal{N} \models T$ that is an extension of $\mathcal{M}$, we have $\psi^{\mathcal{N}}(a_1, \ldots, a_m) = \psi^{\mathcal{M}}(a_1, \ldots, a_m)$.

An L-theory T is *model complete* if any embedding between models of T is an elementary embedding.

3.5 Proposition *The L-theory T is model complete if and only if every model of T is an existentially closed model of T.*

Proof This is Robinson's Criterion for model completeness. The proof given by Hodges [14, Theorem 8.3.1] for classical first order logic can easily be adapted to the continuous setting. □

In Ben Yaacov [2, Appendix A] there is some further discussion of inf- and sup-formulas and of model completeness.

3.6 Definition Let T be an L-theory. A *model companion* of T is an L-theory S such that:

- Every model of S embeds in a model of T.
- Every model of T embeds in a model of S.
- S is model complete.

Note that the first two criteria in this definition together are equivalent to the statement $S_\forall = T_\forall$. As in classical first order logic, if a theory has a model companion, then that model companion is unique (up to equivalence of theories).

Recall that a theory T is *inductive* if whenever Λ is a linearly ordered set and $(\mathcal{M}_\lambda \mid \lambda \in \Lambda)$ is a chain of models of T, then the completion of the union of $(\mathcal{M}_\lambda \mid \lambda \in \Lambda)$ is a model of T.

3.7 Proposition *Let T be an inductive L-theory and let $\mathcal{K}$ be the class of existentially closed models of T. If there exists an L-theory S so that $\mathcal{K} = \mathrm{Mod}(S)$, then S is the model companion of T.*

Proof The proof of [14, Theorem 8.3.6] can be adapted to the continuous setting. □

We say the L-theory T has *amalgamation over substructures* if for any substructures $\mathcal{M}_0$, $\mathcal{M}_1$ and $\mathcal{M}_2$ of models of T and embeddings $f_1\colon \mathcal{M}_0 \to \mathcal{M}_1, f_2\colon \mathcal{M}_0 \to \mathcal{M}_2$, there exists a model $\mathcal{N}$ of T and embeddings $g_i\colon \mathcal{M}_i \to \mathcal{N}$ such that $g_1 \circ f_1 = g_2 \circ f_2$.

3.8 Proposition *Let T_1 and T_2 be L-theories such that T_2 is the model companion of T_1. Assume T_1 has amalgamation over substructures. Then T_2 has quantifier elimination.*

Proof The corresponding result in classical first order logic is the equivalence of (a) and (d) in [14, Theorem 8.4.1]. The proof given there can be adapted to the continuous setting. □

4   The theory of pointed $\mathbb{R}$-trees with radius at most r

By the *radius* of a pointed metric space (M, p) we mean the supremum of the distance $d(p, x)$ as x varies over M, and when using the term we always require that this expression is finite.

In this section we first present the continuous signature used in this paper to study $\mathbb{R}$-trees. We then give axioms for the theory $\mathbb{R}\mathrm{T}_r$ of $\mathbb{R}$-trees with radius $\leq r$.

Let $r > 0$ be a real number. Define the signature $L_r := \{p\}$ where p is a constant symbol and specify that the metric symbol d has values which lie in the interval $[0, 2r]$. Any pointed metric space (M, p) with radius $\leq r$ naturally gives rise to an L_r-prestructure $\mathcal{M} = (M, d, p)$, in which d is a metric; $\mathcal{M}$ is an L_r-structure if the metric space involved is metrically complete.

4.1 Remark For future reference, we note: if M is a pointed $\mathbb{R}$-tree of radius at most r and if $x \in M$ satisfies $d(p, x) = r$, then x is an endpoint in M. (See Definition 2.16.) Indeed, if x is not an endpoint then we can let $y \in M$ be an element of a branch at x that does not contain p. But then we would have $d(p, y) = d(p, x) + d(x, y) > r$.

Next we define a set of axioms $\mathbb{R}\mathrm{T}_r$ for the class of complete, pointed $\mathbb{R}$-trees of radius $\leq r$. Recall the connective $\dot{-} : [0, \infty) \times [0, \infty) \to [0, \infty)$ defined by $x \dot{-} y = \max\{x - y, 0\}$.

4.2 Definition Let $\mathbb{R}\mathrm{T}_r$ be the L_r-theory consisting of the following conditions:

(1)
$$\sup_x d(x, p) \leq r$$

(2)
$$\sup_x \sup_y \inf_z \max\{|d(x, z) - \tfrac{1}{2}d(x, y)|, \ |d(y, z) - \tfrac{1}{2}d(x, y)|\} = 0$$

(3)
$$\sup_x \sup_y \sup_z \sup_w \left(\min\{(x \cdot z)_w, (y \cdot z)_w\} \dot{-} (x \cdot y)_w \right) = 0$$

Condition (2) formalizes the approximate midpoint property. In reading (3), recall that $(x \cdot y)_w$ denotes the Gromov product (see Definition 2.8), which is given by an explicit formula $\varphi(x, y, w)$ in the signature L_r.

The next lemma shows that the class of complete pointed $\mathbb{R}$-trees of radius $\leq r$ is axiomatized by $\mathbb{R}\mathrm{T}_r$.

4.3 Lemma *The models of $\mathbb{R}\mathrm{T}_r$ are exactly the complete, pointed $\mathbb{R}$-trees of radius $\leq r$.*

Proof First we assume $\mathcal{M} \models \mathbb{R}\mathrm{T}_r$. Then (M, d, p) is a complete, pointed metric space. Axiom (1) guarantees that M has radius $\leq r$. Axiom 3 implies that M is 0-hyperbolic.

Axiom (2) implies that for any $x, y \in M$ and any $\epsilon > 0$ there is $z \in M$ such that $d(x, z)$ and $d(y, z)$ are within ϵ of $d(x, y)/2$. We show that in a complete metric space that is 0-hyperbolic, this implies the exact midpoint property. (See Fact 2.2.) Given $x, y \in M$, for each n let $z_n \in M$ be such that $d(x, z)$ and $d(y, z)$ are within $1/n$ of $d(x, y)/2$. Applying the 4-point condition (see Lemma 2.10) we have

$$d(z_m, z_n) + d(x, y) \leq \max \left[d(z_m, x) + d(z_n, y), d(z_m, y) + d(z_n, x) \right]$$
$$\leq \max \left[d(x, y) + \frac{1}{m} + \frac{1}{n}, d(x, y) + \frac{1}{m} + \frac{1}{n} \right] = d(x, y) + \frac{1}{m} + \frac{1}{n}$$

from which we get $d(z_m, z_n) \leq 1/m + 1/n$ for all m, n. Therefore (z_n) converges in M to an exact midpoint between x and y.

Therefore, by Lemma 2.14(1), $\mathcal{M}$ is a pointed $\mathbb{R}$-tree with radius $\leq r$.

That a complete, pointed $\mathbb{R}$-tree with radius $\leq r$ is a model of $\mathbb{R}\mathrm{T}_r$ is clear. $\square$

4.4 Remark Structures in continuous logic are required to be metrically complete, while in general, $\mathbb{R}$-trees are not complete. A pointed $\mathbb{R}$-tree M with radius $\leq r$ can naturally be viewed as an L_r-prestructure, which is an L_r-structure iff it is complete (since the pseudometric on the prestructure is actually a metric). If M is not complete, then its metric completion is known to be an $\mathbb{R}$-tree (see Chiswell [8, Lemma 2.4.14]). Further, the completion of a prestructure is known to be an elementary extension, and therefore the prestructure and its completion are completely equivalent from a model-theoretic perspective. (See Ben Yaacov, Berenstein, Henson and Usvyatsov [4, pages 15–17].) Note that this also means any two pointed $\mathbb{R}$-trees of radius $\leq r$ that have the same metric completion are indistinguishable from a model-theoretic perspective, and that a metrically complete $\mathbb{R}$-tree can be identified model-theoretically with any of its dense sub-prestructures. (However, those metric sub-prestructures are not necessarily $\mathbb{R}$-trees. In particular, they are not necessarily geodesic spaces.)

We close this section by noting a property of $\mathbb{R}\mathrm{T}_r$ that will be used later.

4.5 Lemma *The theory $\mathbb{R}\mathrm{T}_r$ is inductive. That is, the completion of the union of an arbitrary chain of models of $\mathbb{R}\mathrm{T}_r$ is a model of $\mathbb{R}\mathrm{T}_r$.*

Proof The proof of Chiswell [8, Lemma 2.1.14] can be modified to show that the union of an arbitrary chain of pointed $\mathbb{R}$-trees is again a pointed $\mathbb{R}$-tree. Also, the

completion of an $\mathbb{R}$-tree is an $\mathbb{R}$-tree. (See [8, Lemma 2.4.14].) Since base points are preserved by embeddings of models, the radius of the underlying pointed $\mathbb{R}$-tree for the union of a chain is most r.

Alternatively, note that $\mathbb{R}\mathrm{T}_r$ is an $\forall\exists$-theory, and therefore the class of its models is closed under completions of unions of chains. □

5   Some definability in $\mathbb{R}\mathrm{T}_r$

We now discuss the notion of *definability* for subsets of and functions on the underlying $\mathbb{R}$-tree M of a model $\mathcal{M}$ of $\mathbb{R}\mathrm{T}_r$. For background on definable predicates, sets and functions see Ben Yaacov, Berenstein, Henson and Usvyatsov [4]. The first result shows that every closed ball centered at the base point is uniformly quantifier-free 0-definable in models of $\mathbb{R}\mathrm{T}_r$.

5.1 Lemma *Let $s \in [0, r]$. Let $\varphi(x)$ be the quantifier-free formula $d(x, p) \doteq s$. Suppose $\mathcal{M} \models \mathbb{R}\mathrm{T}_r$ and (M, d, p) is the underlying $\mathbb{R}$-tree of $\mathcal{M}$. Evidently the zeroset of $\varphi(x)^{\mathcal{M}}$ is equal to the closed ball $B_s(p)$ of radius s centered at p in M. Then for all $x \in M$ we have $\mathrm{dist}(x, B_s(p)) = \varphi(x)^{\mathcal{M}}$.*

Therefore, the closed ball $B_s(p) \subseteq M$ is uniformly 0-definable with respect to the theory $\mathbb{R}\mathrm{T}_r$.

Proof   It suffices to show that $\varphi^{\mathcal{M}}(x)$ is equal to the distance function $\mathrm{dist}(x, B_s(p))$. For $x \in M$ we know $\varphi^{\mathcal{M}}(x) = 0$ if and only if $d(x, p) \leq s$, ie, if and only if $x \in B_s(p)$. Now, let $x \notin B_s(p)$. Then $\varphi^{\mathcal{M}}(x) = d(x, p) - s$. Let γ be a geodesic segment from p to x with $\gamma(0) = p$. Then:

$$\mathrm{dist}(x, B_s(p)) \leq d(x, \gamma(s)) = d(x, p) - d(p, \gamma(s)) = d(x, p) - s = \varphi^{\mathcal{M}}(x)$$

Now toward a contradiction, assume $d(x, p) - s > \mathrm{dist}(x, B_s(p))$. Then there exists a point in $c \in B_s(p)$ with $d(c, x) < d(x, p) - s$. This implies

$$d(x, p) \leq d(p, c) + d(c, x) \leq s + d(c, x) < s + d(x, p) - s = d(x, p)$$

so $d(x, p) < d(x, p)$, which is a contradiction. Therefore, $\mathrm{dist}(x, B_s(p)) = \varphi(x)^{\mathcal{M}}$. □

Note that the preceding argument only needs that the underlying metric space M is a geodesic space.

Next, we present a short discussion of some specific definable functions, points and sets in models of $\mathbb{R}\mathrm{T}_r$.

Let M be an $\mathbb{R}$-tree and for $s \in [0,1]$ define the function $\nu_s \colon M \times M \to M$ by: $\nu_s(x_1, x_2) =$ the point in $[x_1, x_2]$ with distance $s d(x_1, x_2)$ from x_1 and distance $(1 - s)d(x_1, x_2)$ from x_2.

5.2 Lemma *Let $s \in [0,1]$. The function $\nu_s \colon M \times M \to M$ is uniformly 0-definable in models $\mathcal{M}$ of $\mathbb{R}\mathrm{T}_r$ via a quantifier-free formula.*

Proof Let ψ be the formula:

$$\max\{d(x_1, y) \dotminus sd(x_1, x_2), d(x_2, y) \dotminus (1 - s)d(x_1, x_2)\}$$

Let $\mathcal{M} \models \mathbb{R}\mathrm{T}_r$. In $\mathcal{M}$, the distance $d(\nu_s(x_1, x_2), y)$ is equal to $\psi^{\mathcal{M}}(x_1, x_2, y)$. So the function ν_s is 0-definable via this quantifier-free formula in any model of $\mathbb{R}\mathrm{T}_r$. □

In Definition 2.15 we defined $Y = Y(a, b, c)$, the unique point so that $\{Y\} = [a, b] \cap [b, c] \cap [a, c]$. Next we show that Y is a definable function.

5.3 Theorem *The function $Y \colon M^3 \to M$ that for inputs $x_1, x_2, x_3 \in M$ returns $Y(x_1, x_2, x_3)$ is uniformly 0-definable in models $\mathcal{M}$ of $\mathbb{R}\mathrm{T}_r$ via a quantifier-free formula.*

Proof Recall that in an $\mathbb{R}$-tree, the Gromov product $(a \cdot b)_x$ is equal to the distance from x to the segment $[a, b]$, and this distance is realized by a unique closest point on $[a, b]$. Let

$$\varphi(x_1, x_2, x_3, x) = \max\left\{(x_1 \cdot x_2)_x, (x_1 \cdot x_3)_x, (x_2 \cdot x_3)_x\right\}.$$

Let $\mathcal{M} \models \mathbb{R}\mathrm{T}_r$ and $x_1, x_2, x_3 \in M$. We will show that in $\mathcal{M}$,

$$d(x, Y(x_1, x_2, x_3)) = \varphi^{\mathcal{M}}(x_1, x_2, x_3, x).$$

Let E be the closed subtree of M spanned by x_1, x_2 and x_3. Abbreviate $Y(x_1, x_2, x_3)$ by Y. For $x \in M$, let z be the unique point closest to x in E. Then we have $d(x, Y) = d(x, z) + d(z, Y)$. The point z must lie on at least one of $[x_1, Y]$, $[x_2, Y]$ or $[x_3, Y]$. Without loss of generality, assume $z \in [x_1, Y] = [x_1, x_3] \cap [x_1, x_2]$.

Since z is closest in E to x and $z \in [x_1, x_3] \cap [x_1, x_2]$ we know $d(x, z) = (x_1 \cdot x_2)_x = (x_1 \cdot x_3)_x$. That z is closest in E to x also implies $(x_2 \cdot x_3)_x \geq d(x, z)$. This makes $(x_2 \cdot x_3)_x \geq (x_1 \cdot x_2)_x = (x_1 \cdot x_3)_x$. Thus, $\varphi^{\mathcal{M}}(x_1, x_2, x_3, x) = (x_2 \cdot x_3)_x$.

Since $z \in [x_1, Y]$, the point Y is the closest point to z on $[x_2, x_3]$, and therefore Y is the closest point to x on $[x_2, x_3]$. It follows that, $d(x, Y) = (x_2 \cdot x_3)_x$. We conclude that $\varphi^{\mathcal{M}}(x_1, x_2, x_3, x) = d(x, Y)$. □

6 Amalgamation

Next, we discuss amalgamation for the L_r-theory $\mathbb{R}T_r$. The following result from Chiswell [8] discussing amalgamating over points in $\mathbb{R}$-trees is then extended to amalgamation of $\mathbb{R}$-trees over subtrees. This leads to a proof of amalgamation over substructures for $\mathbb{R}T_r$.

6.1 Lemma *Let (X, d) be an $\mathbb{R}$-tree and let $\{(X_i, d_i) \mid i \in I\}$ be a family of $\mathbb{R}$-trees such that $X_i \cap X = \{x_i\}$ for all $i \in I$. $X_i \cap X_j = \{x_i\}$ if $x_i = x_j$ and $X_i \cap X_j = \emptyset$ otherwise. Define $N = (\bigcup_{i \in I} X_i) \cup X$, and define $\hat{d} \colon N \times N \to \mathbb{R}$ by:*

- *On $X \times X$, $\hat{d} = d$, and on $X_i \times X_i$, $\hat{d} = d_i$.*
- *If $x \in X_i$ and $x' \in X_j$, with $i \neq j$, then*

$$\hat{d}(x, x') = \hat{d}(x', x) = d_i(x, x_i) + d(x_i, x_j) + d_j(x_j, x').$$

- *If $x \in X_i$ and $x' \in X$ then $\hat{d}(x, x') = \hat{d}(x', x) = d_i(x, x_i) + d(x_i, x')$.*

Then $(N, \hat{d})$ is an $\mathbb{R}$-tree.

Proof See [8, Lemma 2.1.13] □

We next apply this lemma to prove $\mathbb{R}$-trees can be amalgamated over subtrees.

6.2 Lemma *Given $\mathbb{R}$-trees M_0, M_1 and M_2 and isometric embeddings $f_i \colon M_0 \to M_i$ for $i = 1, 2$, there exists an $\mathbb{R}$-tree N and isometric embeddings $g_i \colon M_i \to N$ such that $g_1 \circ f_1 = g_2 \circ f_2$.*

Proof Without loss of generality, we may assume that M_0 is nonempty, that $M_0 \subseteq M_1$ and $M_0 \subseteq M_2$, that $M_1 \cap M_2 = M_0$, and that the f_i are inclusion maps. Further, we may assume that M_i is complete for $i = 0, 1, 2$, since the maps f_1, f_2 are isometric, and can thus be extended as needed. (Recall that the completion of an $\mathbb{R}$-tree is again an $\mathbb{R}$-tree.) Let $\{x_k \mid k \in |M_0|\}$ be a $1 - 1$ list of the elements of M_0. For each $k \in |M_0|$ let B_k be the set of branches β at x_k in M_2 such that $\beta \cap M_0 = \emptyset$. Let α_k be the cardinality of B_k (which we may assume is nonempty, by extending M_2 if needed), and let $A_k = \{\beta_{kv} \mid v \in \alpha_k\}$ be a $1 - 1$ list of the distinct elements of B_k. Then let $I = \{(k, v) \mid k \in |M_0|, v \in \alpha_k\}$, and define $X_i = \beta_{kv} \cup \{x_k\}$ for each $i = (k, v) \in I$. For each $i = (k, v) \in I$, let $x_i = x_k$. For each $i \in I$ we consider X_i as a subtree of M_2.

We claim that $X = M_1$ and the family $(X_i \mid i \in I)$ satisfy the hypotheses of Lemma 6.1. By construction, every point in M_0 is contained in some X_i. For each $i = (k, v) \in I$

we see that $X_i \cap M_1 = \{x_k\} = \{x_i\}$ since x_k is the only point in $X_i \cap M_0$ and $X_i \cap M_1 \subseteq M_2 \cap M_1 = M_0$. Further, given any point $y \in M_2 \setminus M_0$, we may let x be the closest point to y in M_0 and take k so that $x_k = x$; then we may take β to be a branch in M_2 at x_k that contains y. It follows that $\beta \cap M_0 = \emptyset$; indeed, if $x' \in \beta \cap M_0$, then $Y(x, x', y)$ would be an element of M_0 that was closer to y than x. Thus there exists $i = (k, v) \in I$ for which $\beta = \beta_i$ and hence $y \in X_i$. Thus we have $M_2 = \bigcup_{i \in I} X_i$.

Applying Lemma 6.1, we get that $N = (\bigcup_{i \in I} X_i) \cup M_1$ with the metric $\hat{d}$ is an $\mathbb{R}$-tree, and clearly $N = M_2 \cup M_1$ as metric spaces. Define g_i to be the inclusion of M_i in N. Then $g_1 \circ f_1 = g_2 \circ f_2$ is clear. $\qquad\square$

To apply these results to models of $\mathbb{R}\mathrm{T}_r$, we begin by proving any subset of a model of $\mathbb{R}\mathrm{T}_r$ gives rise to a unique model of $\mathbb{R}\mathrm{T}_r$. In particular, any substructure extends to a unique model.

6.3 Lemma *Let $\mathcal{N} \models \mathbb{R}\mathrm{T}_r$. Any subset $A \subseteq N$ extends to a model $\mathcal{M} \models \mathbb{R}\mathrm{T}_r$ such that whenever f is an embedding of $A \cup \{p\}$ into $\mathcal{W} \models \mathbb{R}\mathrm{T}_r$ that satisfies $f(p) = p$, there is a unique extension of f that embeds $\mathcal{M}$ into $\mathcal{W}$.*

Proof Let (N, d, p) be the underlying $\mathbb{R}$-tree of the model $\mathcal{N}$, and let $A \subseteq N$. Define the structure $\mathcal{M} = (M, d, p)$ by $M = \overline{E_{A \cup \{p\}}}$ with the metric and base point from (N, d, p). Then M is the smallest closed subtree of N containing $A \cup \{p\}$. Obviously $\mathcal{M} \models \mathbb{R}\mathrm{T}_r$; Lemma 2.14(3) yields a unique extension of f to embed $E_{A \cup \{p\}}$ into $\mathcal{W}$. Further, since f is an isometry, it extends further (also uniquely) to embed $\mathcal{M}$ into $\mathcal{W}$. $\qquad\square$

6.4 Theorem *The L_r-theory $\mathbb{R}\mathrm{T}_r$ has amalgamation over substructures. That is, if $\mathcal{M}_0$, $\mathcal{M}_1$ and $\mathcal{M}_2$ are substructures of models of $\mathbb{R}\mathrm{T}_r$ and $f_1 \colon \mathcal{M}_0 \to \mathcal{M}_1$, $f_2 \colon \mathcal{M}_0 \to \mathcal{M}_2$ are embeddings, then there exists a model $\mathcal{N}$ of $\mathbb{R}\mathrm{T}_r$ and embeddings $g_i \colon \mathcal{M}_i \to \mathcal{N}$ such that $g_1 \circ f_1 = g_2 \circ f_2$.*

Proof First, by Lemma 6.3 we may assume that $\mathcal{M}_0$, $\mathcal{M}_1$ and $\mathcal{M}_2$ are models of $\mathbb{R}\mathrm{T}_r$ with underlying complete, pointed $\mathbb{R}$-trees (M_i, d_i, p_i). Further, we may assume that $M_0 \subseteq M_i$ and $p_0 = p_i$ for $i = 1, 2$ and that the f_i are inclusion maps.

We use Lemma 6.2 to construct the $\mathbb{R}$-tree (N, d) and isometric embeddings $g_i \colon M_i \to N$ such that $g_1 \circ f_1 = g_2 \circ f_2$. Recall that in this construction, $N = M_1 \cup M_2$. Define the base point of N to be $q = p_1 = p_2 = p_0$. Since every point in each M_i has distance $\leq r$ from $p_i = q$, we conclude that the pointed $\mathbb{R}$-tree (N, d, q) has radius

$\leq r$. Let $\mathcal{N} = (N, d, q)$ be the corresponding L_r-structure. Then $\mathcal{N}$ is a model of $\mathbb{R}\mathrm{T}_r$, and since g_i are isometric embeddings preserving the base point, they give rise to embeddings of L_r-structures as required.   □

7   The model companion of $\mathbb{R}\mathrm{T}_r$

In this section we define what it means for a pointed $\mathbb{R}$-tree of radius $\leq r$ to be *richly branching*. We then show that the theory of richly branching pointed $\mathbb{R}$-trees with radius r is the model companion of $\mathbb{R}\mathrm{T}_r$. Throughout the rest of this paper we assume $r > 0$.

7.1 Definition   A pointed $\mathbb{R}$-tree (M, d, p) of radius $\leq r$ is *richly branching* if the set

$$B = \{b \in M \mid \text{at } b \text{ there are at least 3 branches of height } \geq r - d(p, b)\}$$

is dense in M.

7.2 Remark   The preceding definition and arguments below apply specifically to pointed $\mathbb{R}$-trees with radius $\leq r$. For general $\mathbb{R}$-trees we define: *an $\mathbb{R}$-tree M is richly branching if the set of points at which there are at least 3 branches of infinite height is dense in M.* Note that an $\mathbb{R}$-tree that is richly branching in this sense must be unbounded.

7.3 Lemma   *Let (X, d) be an $\mathbb{R}$-tree in which the set of branch points is dense. Then for any distinct $x, y \in X$, the set of branch points on $[x, y]$ is dense in $[x, y]$.*

Proof   Let z be on the segment $[x, y]$ and suppose $0 < \delta < \min[d(x, z), d(y, z)]$. By assumption there is a branch point w in X with $d(w, z) < \delta$. If w lies on $[x, y]$ we are done. Otherwise, let u be the point on $[x, y]$ that is closest to w. By Lemma 2.5 we have that x, y, w are in distinct branches at u, and $d(z, u) \leq d(w, u) + d(z, u) = d(w, z) < \delta$. That is, u is a branch point on $[x, y]$ that is arbitrarily close to an arbitrary point on $[x, y]$ that is distinct from x, y.   □

If we start with an unbounded richly branching $\mathbb{R}$-tree M, assign an arbitrary base point p and select $r > 0$, we can make a richly branching $\mathbb{R}$-tree with radius $\leq r$ as in Definition 7.1 by taking the closed ball of radius r with center p in M. Note that this yields an L_r-prestructure and the completion is an L_r-structure.

Conversely, suppose M is an $\mathbb{R}$-tree and p is any point in M. If the closed ball of radius r with center p in M is richly branching in the sense of 7.1 for an unbounded set of $r > 0$, then M is richly branching as a general $\mathbb{R}$-tree. (See Lemma 7.7 below.)

Next we give axioms for the class of complete, richly branching, pointed $\mathbb{R}$-trees with radius r.

7.4 Definition   Define $\psi(x)$ to be the L_r-formula

$$\inf_{y_1 y_2 y_3} \max \left\{ \max_{i=1,2,3} \{|d(x,y_i) - (r - d(p,x))|\}, \max_{1 \le i < j \le 3} \{d(x,y_i) + d(x,y_j) - d(y_i,y_j)\} \right\}$$

and let $\varphi = \sup_x \psi(x)$.

7.5 Definition   Let $\mathrm{rb}\mathbb{R}\mathrm{T}_r = \mathbb{R}\mathrm{T}_r \cup \{\varphi = 0\}$.

To help parse these axioms and picture what they mean, note that in an $\mathbb{R}$-tree:

$$d(x,y_i) + d(x,y_j) - d(y_i,y_j) = 2(y_i \cdot y_j)_x = 2\,\mathrm{dist}(x, [y_i, y_j])$$

Consider the case when the infima are realized exactly in $\mathcal{M}$ (eg, if $\mathcal{M}$ is an ω-saturated model.) In this setting, $\varphi = 0$ means that for any element $a \in M$ with $d(p,a) < r$, there exist b_1, b_2, b_3 each on a separate branch at a and each with distance $r - d(p,a)$ from a. We know b_1, b_2, b_3 are on distinct branches at a because $d(a,b_i) > 0$ for every $1 \le i \le 3$ and $2\,\mathrm{dist}(a, [b_i, b_j]) = 0$ for every $1 \le i < j \le 3$. In particular, this makes a a branch point. In a general model where the infima are not necessarily realized, an approximate version of this is true; given an element a there must be a branch point within ϵ of a, as we now show.

7.6 Lemma   *Let $\mathcal{M} \models \mathrm{rb}\mathbb{R}\mathrm{T}_r$ with underlying pointed $\mathbb{R}$-tree (M, d, p) and let $a \in M$. Let $h \in \mathbb{R}$ be such that $0 < h < r - d(p,a)$. For any $\epsilon > 0$ there exists a point $b \in M$ so that $d(a,b) < \epsilon$ and there are 3 branches at b, each with height at least h.*

Proof   Let $s = r - d(p,a)$, and assume $0 < 4\epsilon < s - h$. Since $\varphi^{\mathcal{M}} = 0$, we know $\psi(a)^{\mathcal{M}} = 0$. Thus there exist points $c_1, c_2, c_3 \in M$ so that

$$|d(a,c_i) - s| < \epsilon \quad \text{and} \quad d(a,c_i) + d(a,c_j) - d(c_i,c_j) < \epsilon.$$

It follows that for $i \ne j$

$$d(c_i,c_j) > d(a,c_i) + d(a,c_j) - \epsilon > 2s - 3\epsilon > 0$$

meaning c_1, c_2, c_3 must be distinct. Moreover, c_1, c_2 and c_3 cannot lie along a single piecewise segment. To see this assume $[c_i, c_k] = [c_i, c_j, c_k]$ for $i, j, k \in \{1, 2, 3\}$. Then $d(c_i, c_k) = d(c_i, c_j) + d(c_j, c_k)$, implying $d(c_i, c_k) = d(c_i, c_j) + d(c_j, c_k) > 4s - 6\epsilon$. However, $d(c_i, c_k) \le d(a, c_i) + d(a, c_k) < 2s + 2\epsilon$. These inequalities imply $4\epsilon > s >$

$s - h$, a contradiction. So, each of c_1, c_2, c_3 lies on a different branch at $Y(c_1, c_2, c_3)$. Thus, there are at least 3 branches at $Y(c_1, c_2, c_3)$.

Next, let $u = Y(a, c_1, c_2)$, $v = Y(a, c_2, c_3)$ and $w = Y(a, c_1, c_3)$. Without loss of generality, we assume $d(a, u) \leq d(a, v) \leq d(a, w)$; in other words, $\mathrm{dist}(a, [c_1, c_2]) \leq \mathrm{dist}(a, [c_2, c_3]) \leq \mathrm{dist}(a, [c_1, c_3])$. Then $\mathrm{dist}(a, [c_1, c_2]) = \mathrm{dist}(a, [c_2, c_3])$, since M is 0-hyperbolic. Lemma 2.1.6 in Chiswell [8] and the subsequent discussion yield that $u = v$, so the point closest to a on $[c_1, c_2]$ and the point closest to a on $[c_2, c_3]$ are the same and $w = Y(c_1, c_2, c_3)$. By Lemma 2.4, we have $2\,\mathrm{dist}(a, [c_1, c_3]) = d(a, c_1) + d(a, c_3) - d(c_1, c_3)$. Thus, $d(a, w) = \mathrm{dist}(a, [c_1, c_3]) < \epsilon/2 < \epsilon$. So, $w = Y(c_1, c_2, c_3)$ has distance $< \epsilon$ from a. Finally, for $i = 1, 2, 3$ we have:

$$d(w, c_i) \geq d(a, c_i) - d(a, w) > s - \epsilon - \epsilon/2 > s - 4\epsilon > h$$

Therefore $b = Y(c_1, c_2, c_3)$ satisfies the desired conditions. $\qquad\qquad\qquad\qquad$ □

The next lemma shows that any branch at any point in a model of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ must have maximum possible height, with that height realized by a point having distance r from p. This also implies every model of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ has radius equal to r.

7.7 Lemma *Let $\mathcal{M} \models \mathrm{rb}\mathbb{R}\mathrm{T}_r$ and let $a \in M$. In any branch β at a, there exists at least one point b such that $d(p, b) = r$.*

Proof Let $a \in M$ and $\delta = r - d(p, a)$, and let β be a branch at a.

We first assume $p \notin \beta$ (including the case where $a = p$). Thus a is not an endpoint in M, so we have $\delta > 0$ by Remark 4.1. In what follows, by iterating the use of Lemma 7.6, we build a sequence $b_1, b_2, b_3, \ldots$ of points in β so that for any $k \in \mathbb{N}$, $[p, a, b_1, \ldots, b_k]$ is a piecewise segment, and $|r - d(p, b_i)| = r - d(p, b_i) < \delta/2^i$.

First we choose b_1. Let $c \in \beta$, so $d(a, c) > 0$. If c is such that $r - d(p, c) < \delta/2$, then let $b_1 = c$. Since c and p are in different branches at a, it follows that $d(a, c) \leq \delta$. If $r - d(p, c) \geq \delta/2$, then let $0 < \epsilon < d(a, c)/2$. Use Lemma 7.6 to find c' so that $d(c, c') < \epsilon$, and so that there are at least 3 branches at c' with height at least $(1 - d(a, c)/(2\delta))(r - d(p, c))$. So $|d(a, c) - d(a, c')| < \epsilon$, and since $\epsilon < d(a, c)/2$, we know $c' \in \beta$. Then find a point b_1 on a branch at c' other than the one containing a. This makes $[p, a, c', b_1]$ a piecewise segment. Select b_1 so that:

$$d(c', b_1) = (1 - \frac{d(a, c)}{\delta})(r - d(p, c)) \geq (1 - \frac{d(a, c)}{\delta})\frac{\delta}{2} = \frac{\delta}{2} - \frac{d(a, c)}{2}$$

Then: $\quad d(a, b_1) = d(a, c') + d(c', b_1) > (d(a, c) - \epsilon) + d(c', b_1)$

$$> d(a, c) - \frac{d(a, c)}{2} + \frac{\delta}{2} - \frac{d(a, c)}{2} = \frac{\delta}{2}$$

Thus $d(a, b_1) > \delta/2$. It follows that:

$$r - d(p, b_1) = r - (d(p, a) + d(a, b_1)) = (r - d(p, a)) - d(a, b_1)$$

$$= \delta - d(a, b_1) < \delta - \frac{\delta}{2} = \frac{\delta}{2}$$

Once we have $b_1, \ldots, b_i$, we proceed in a manner analogous to what was done above to find b_{i+1} so that $|r - d(p, b_i)| = r - d(p, b_i) < \delta/2^i$. Then for any $j > i$ we have:

$$d(b_i, b_j) = d(p, b_j) - d(p, b_i) < r - (r - \frac{\delta}{2^i}) = \frac{\delta}{2^i}$$

It follows that (b_n) is a Cauchy sequence. Let b be the limit of (b_i).

Claim: $b_i \in [p, b]$ for all $i \in \mathbb{N}$.

If not, let $i \in \mathbb{N}$ be such that $b_i \notin [p, b]$. Then for all $j \geq i$, $b_j \notin [p, b]$, because $[p, b_i, b_j]$ is a piecewise segment. Let $j > i$ and let q be the closest point to b on $[p, b_i, b_j]$, ie, $q = Y(p, b_j, b)$. Either $q \in [p, b_i]$ or $q \in [b_i, b_j]$. If $q \in [b_i, b_j]$, then $[p, q] = [p, b_i, q]$ implying $b_i \in [p, b] = [p, b_i, q, b]$, a contradiction. Thus, $q \in [p, b_i]$, and $[p, q, b_i, b_j]$ is a piecewise segment. Using these facts and our choice of q we conclude:

$$d(b, b_j) = d(b, q) + d(q, b_j) = d(b, q) + d(q, b_i) + d(b_i, b_j)$$

$$\geq d(b, q) + d(q, b_i) = d(b, b_i)$$

We have demonstrated that $d(b, b_j) \geq d(b, b_i) > 0$ for all $j > i$, contradicting that b is the limit of the sequence (b_i). This proves the claim.

Because $b_i \in [p, b]$, we know $r \geq d(p, b) \geq d(p, b_i) > r - \delta/2^i$ for any $i \in \mathbb{N}$. Therefore, $d(p, b) = r$. And $b \in \beta$, since otherwise b would be on a different branch at a than b_1. That would make $[b_1, a, b]$ a piecewise segment, contradicting that $[p, b] = [p, a, b_1, b]$ is a piecewise segment.

Finally, we deal with the case where $p \in \beta$. Since β is open, we may use Lemma 7.6 to find a point $a' \in \beta$ such that there are 3 branches at a'. Select a branch β' at a' so that $p \notin \beta'$ and $a \notin \beta'$. The latter guarantees $\beta' \subseteq \beta$. Applying the first part of this proof to a' and β' yields a point $b \in \beta'$ such that $d(p, b) = r$. $\qquad\square$

The following theorem shows that complete richly branching pointed $\mathbb{R}$-trees with radius r form an elementary class.

7.8 Theorem *The models of* rb$\mathbb{R}$T$_r$ *are exactly the complete, richly branching* $\mathbb{R}$*-trees with radius* r.

Proof Let (M, d, p) be a complete, richly branching pointed $\mathbb{R}$-tree with radius r and let $\mathcal{M}$ be the corresponding L_r-structure. Clearly $\mathcal{M} \models \mathbb{R}\mathrm{T}_r$, and it remains to verify that $\varphi^{\mathcal{M}} = 0$. Let $a \in M$. If $d(p, a) = r$, then let $c_1 = c_2 = c_3 = a$ and note that these witness $\psi^{\mathcal{M}}(a) = 0$. So we assume $d(p, a) < r$, and let $\epsilon > 0$ be such that $\epsilon/3 < r - d(p, a)$. Since M is richly branching, the set B from Definition 7.1 is dense in M. So, there exists $b \in B$ with $d(a, b) < \epsilon/3$. Then there exist $c_1, c_2, c_3 \in M$ such that each distance $d(b, c_i)$ is $> r - d(b, p) - \epsilon/3$, and each of c_1, c_2, c_3 is on a different branch at b. By moving these points closer to b along the geodesic segments if necessary, we may assume $|d(b, c_i) - (r - d(b, p))| < \epsilon/3$ for each $i = 1, 2, 3$. The triangle inequality in M implies $|d(a, c_i) - d(b, c_i)| < \epsilon/3$ for all $i = 1, 2, 3$, and $|(r - d(p, a)) - (r - d(p, b))| < \epsilon/3$. The triangle inequality in $\mathbb{R}$ yields $|d(a, c_i) - (r - d(p, a))| < \epsilon$ for each $i = 1, 2, 3$.

Next we fix $i \neq j$ in $\{1, 2, 3\}$ and note that $d(c_i, c_j) = d(c_i, b) + d(c_j, b)$ because c_1, c_2, c_3 are on different branches at b. Since we have $d(a, b) < \epsilon/3$, the triangle inequality yields $|d(a, c_i) + d(a, c_j) - d(c_i, c_j)| < 2\epsilon/3$.

So, for each ϵ, there are c_1, c_2 and c_3 making

$$\max_{i=1,2,3} \left\{ |d(a, c_i) - (r - d(p, a))| \right\} < \epsilon$$

and

$$\max_{1 \leq i < j \leq 3} \left\{ |(d(a, c_i) + d(a, c_j)) - d(c_i, c_j)| \right\} < \epsilon.$$

Thus, $\psi(a)^{\mathcal{M}} = 0$, and since a was arbitrary, we conclude $\varphi^{\mathcal{M}} = 0$.

For the converse direction, assume $\mathcal{M} \models \mathrm{rb}\mathbb{R}\mathrm{T}_r$, so (M, d, p) is a complete, pointed $\mathbb{R}$-tree. Let $a \in M$ and $\epsilon > 0$. By Lemma 7.6 we may find $b \in M$ so that $d(a, b) < \epsilon$ and there are at least 3 distinct branches at b. By Lemma 7.7 each branch at b contains a point with distance r from p, so the height of each branch is at least $r - d(p, b)$. It also follows from 7.7 that M has radius r. Since $a \in M$ and $\epsilon > 0$ were arbitrary, we conclude the set B of points with at least 3 branches of height $\geq r - d(p, b)$ is dense in M. Therefore, M is richly branching. $\square$

Next, we turn to a series of lemmas that are needed for our proof that the theory $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ is the model companion of $\mathbb{R}\mathrm{T}_r$. (See Theorem 7.12 below.)

7.9 Lemma *Every existentially closed model of $\mathbb{R}\mathrm{T}_r$ is a model of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$.*

Proof Let $\mathcal{M} \models \mathbb{R}\mathrm{T}_r$ be existentially closed with underlying $\mathbb{R}$-tree (M, d, p). Consider $a \in M$. If $r - d(p, a) = 0$, then $\psi^{\mathcal{M}}(a) = 0$ is true: make $y_1 = y_2 = y_3 = a$.

When $r - d(p, a) > 0$, using Lemma 6.1 we may construct an extension $\mathcal{N} \models \mathbb{R}T_r$ of $\mathcal{M}$ with underlying $\mathbb{R}$-tree (N, d, p) such that in N there exist c_1, c_2 and c_3, each on a different branch at a and each with distance $r - d(p, a)$ from a. It follows that $\psi^{\mathcal{N}}(a) = 0$. Since $\mathcal{M}$ is existentially closed and ψ is an inf-formula, we have $\psi^{\mathcal{N}}(a) = \psi^{\mathcal{M}}(a) = 0$. Our choice of $a \in M$ was arbitrary, so $\varphi^{\mathcal{M}} = 0$. It follows that $\mathcal{M}$ is a model of rb$\mathbb{R}T_r$. □

The next lemma connects κ-saturation to the number of branches at the points whose distance from the base point is $< r$, in a richly branching $\mathbb{R}$-tree of radius r. (Recall from Remark 4.1 that every point at distance $= r$ from the base point in such an $\mathbb{R}$-tree is an endpoint.) The converse of this lemma is also true, and the resulting characterization of κ-saturated models of rb$\mathbb{R}T_r$ is Theorem 8.5 below, which will be proved once we have shown that rb$\mathbb{R}T_r$ admits quantifier elimination.

7.10 Lemma *Let $\mathcal{M} \models$ rb$\mathbb{R}T_r$. Let κ be an infinite cardinal. If $\mathcal{M}$ is κ-saturated, then (M, d, p) has at least κ-many branches at every point a such that $d(p, a) < r$.*

Proof We begin with the case where $\kappa = \omega$. Arguing by contradiction, we assume there exists an ω-saturated model $\mathcal{M}$ of rb$\mathbb{R}T_r$ containing a point a with $d(p, a) < r$ such that the degree of a is finite. Let $\beta_1, \ldots, \beta_n$ be the distinct branches at a.

For each $i = 1, \ldots, n$ we can find $b_i \in \beta_i$ so that $d(a, b_i) = r - d(p, a)$ as follows. When p is not in the branch β_i we let b_i be a point in β_i with $d(p, b_i) = r$. The existence of such a point is guaranteed by Lemma 7.7. In this case, either $p = a$, or p and b_i are on different branches at a, and it follows in both instances that $d(a, b_i) = r - d(p, a)$. When the base point p is in branch β_i, we select $x \in [a, p]$ distinct from a and p such that there are at least 3 branches at x. Then $d(p, x) < d(p, a)$, so $r - d(p, x) > r - d(p, a)$. On a branch at x that contains neither p nor a, by Lemma 7.7, there is a point y so that $d(p, y) = r$. Note that $(a, y] \subseteq \beta_i$. Since p and y are on different branches at x, we know $d(p, x) + d(x, y) = d(p, y) = r$, so $d(x, y) = r - d(p, x)$. Then

$$d(a, y) = d(a, x) + d(x, y) = d(a, x) + (r - d(p, x)) > d(a, x) + (r - d(p, a)) > r - d(p, a)$$

because a and y are on different branches at x. Since $d(a, y) > r - d(p, a)$ we may take the point on $(a, y]$ with distance $r - d(p, a)$ from a to be our b_i.

Next, let $r - d(p, a) > \epsilon > 0$. Since M is a richly branching $\mathbb{R}$-tree, there exists a' in branch β_1 such that $0 < d(a, a') < \epsilon/2$ and there are at least 3 branches at a'. Choose a branch at a' that does not contain a and does not contain b_1. Using Lemma 7.7, find c in this branch at a' so that $d(a', c) = r - d(p, a')$. Note that by Lemma 2.4 we

know $[b_1, a', c]$ and $[a, a', c]$ are both piecewise segments, as is $[b_i, a, a', c]$ for any $i = 2, \ldots, n$.

Using these facts and the fact that $d(a, a') < \epsilon/2$, it is now straightforward to show that

$$|d(a, c) - (r - d(p, a))| < \epsilon$$

and
$$|d(b_i, c) - (d(b_i, a) + d(a, c))| < \epsilon.$$

Let Σ be the set of all conditions

$$|d(a, x) - (r - d(p, a))| \leq 1/k$$

and
$$|d(b_i, x) - (d(b_i, a) + d(a, x))| \leq 1/k$$

for $k \in \mathbb{N}$ and $i = 1, \ldots, n$.

Since ϵ was arbitrary, we have that Σ is finitely satisfiable. Since there are only finitely many parameters in Σ, and $\mathcal{M}$ is ω-saturated, there must exist $b \in M$ that satisfies all of these conditions simultaneously. For this b we have $d(a, b) = r - d(p, a)$ and $d(b_i, b) = d(b_i, a) + d(a, b)$, meaning b is in a different branch at a from any of the b_i. This contradicts that each branch at a was represented by one of $b_1, \ldots, b_n$. Therefore there must be at least ω-many branches at each point in M when $\mathcal{M}$ is ω-saturated.

Now, let $\kappa > \omega$ and assume $\mathcal{M}$ is κ-saturated. Let $a \in M$ with $d(p, a) < r$. Assume toward a contradiction that there are exactly α-many distinct branches at a where $\alpha < \kappa$. Index the branches at a by $i < \alpha$, and by Lemma 7.7 on each of these branches designate a point b_i such that $d(a, b_i) = r - d(p, a)$. Let $A = \{a\} \cup \{b_i \mid i < \alpha\}$ and note that A has cardinality less than κ. Define:

$$\Sigma = \{|d(a, x) - (r - d(p, a))| = 0\} \cup \{|d(b_i, x) - (d(b_i, a) + d(a, x))| = 0 \mid i < \alpha\}$$

It is straightforward to show Σ is finitely satisfiable, since by the first part of this proof there must be infinitely many branches at a. By κ-saturation there exists $b \in M$ that satisfies all of these conditions. Using Lemma 2.4 we conclude that b is on a different branch out of a from each b_i, contradicting the assumption that every branch at a was represented by one of the b_i. $\square$

7.11 Lemma *Let κ be an infinite cardinal. Assume $\mathcal{M}$ is a κ-saturated model of* rb$\mathbb{R}$T$_r$ *with underlying $\mathbb{R}$-tree (M, d, p). Let K be a non-empty, finitely spanned $\mathbb{R}$-tree. Designate a base point q in K. For any $e \in M$ such that $r - d(p, e) \geq \sup\{d(q, x) \mid x \in K\}$ and any collection $\{\beta_i \mid i < \alpha\}$ of branches at e with cardinality $\alpha < \kappa$, there exists an isometric embedding f of K into M such that $f(q) = e$ and $f(K) \cap \beta_i = \emptyset$ for all $i < \alpha$.*

Proof By Lemma 7.10, (M, d, p) has at least κ-many branches at each point a satisfying $d(p, a) < r$. By Lemma 2.17, there is a minimal set that spans K, namely, the set of endpoints of K. Proceed by induction on the size of this minimal spanning set, building up the embedding at each step using the fact that there are κ-many branches of sufficient height at every interior point. The restriction that $r - d(p, e) \geq \sup\{d(q, x) \mid x \in K\}$ keeps the image of the embedding inside M. $\qquad\square$

7.12 Theorem *The L_r-theory* $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ *is the model companion of* $\mathbb{R}\mathrm{T}_r$.

Proof Since $\mathbb{R}\mathrm{T}_r$ is an inductive theory, by Lemma 3.7 it suffices to show that the models of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ are exactly the existentially closed models of $\mathbb{R}\mathrm{T}_r$. By Lemma 7.9 we know every existentially closed model of $\mathbb{R}\mathrm{T}_r$ is a model of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$. It remains to show that every model of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ is an existentially closed model of $\mathbb{R}\mathrm{T}_r$. Let $\mathcal{M} \models \mathrm{rb}\mathbb{R}\mathrm{T}_r$. Let $\mathcal{N} \models \mathbb{R}\mathrm{T}_r$ be an extension of $\mathcal{M}$. We may assume $\mathcal{M}$ and $\mathcal{N}$ are ω_1-saturated. (This is because we may consider the structure $\big(\mathcal{M}, \mathcal{N}, \iota\big)$ where ι is the embedding from $\mathcal{M}$ to $\mathcal{N}$, and take an ω_1-saturated elementary extension of that structure. If we can verify the definition of existentially closed in that setting, it will be true of $\mathcal{M}$ and $\mathcal{N}$.) Let (M, d, p) and (N, d, p) be the underlying pointed $\mathbb{R}$-trees for $\mathcal{M}$ and $\mathcal{N}$ respectively.

Let $a = a_1, \ldots, a_k \in M$, and without loss of generality assume base point p is among $a_1, \ldots, a_k$. We claim that for any $b_1, \ldots, b_l \in N$, there exist $c_1, \ldots, c_l \in M$ so that $d(b_i, b_j) = d(c_i, c_j)$ for all $i, j \in \{1, \ldots, l\}$ and $d(a_i, b_j) = d(a_i, c_j)$ for all $i \in \{1, \ldots, k\}$ and $j \in \{1, \ldots, l\}$.

To prove this claim, let $b_1, \ldots, b_l \in N$, and let $E_a \subseteq M$ be the subtree spanned by $a = a_1, \ldots, a_k$. Note that $E_a = \overline{E_a}$. Define an equivalence relation on $b_1, \ldots, b_l$ by: $b_i \sim b_j$ if b_i and b_j have the same closest point in E_a. Let $A_1, \ldots, A_m$ be the equivalence classes of this equivalence relation, and for $v = 1, \ldots, m$ let e_v be the unique closest point in E_a common to the members of A_v.

For each $v = 1, \ldots, m$ let K_v be the $\mathbb{R}$-tree spanned by $A_v \cup \{e_v\}$ in N, with base point e_v. Note that each K_v is closed and for $u \neq v$, $K_u \cap K_v = \emptyset$. Since (N, d, p) has radius r and $p \in E_a$ and e_v is the unique closest point to x in E_a, Lemma 2.4 gives $d(p, e_v) + d(e_v, x) = d(p, x) \leq r$ for each $x \in K_v$. Therefore $d(e_v, x) \leq r - d(p, e_v)$ for each $x \in K_v$. Thus, $\sup\{d(e_v, x) \mid x \in K_v\} \leq r - d(p, e_v)$. Now, for each $v = 1, \ldots, m$, by Lemma 7.11, there is an isometric embedding $f_v \colon K_v \to M$ sending e_v to e_v such that $f_v(K_v)$ does not intersect E_a except at e_v. Note that e_v is the unique closest point in E_a for every point in $f_v(K_v)$.

Let f be the union of the functions f_v for $v = 1, \ldots, m$. If b_i and b_j are both in A_v, then $d(b_i, b_j) = d(f_v(b_i), f_v(b_j)) = d(f(b_i), f(b_j))$. If b_i and b_j are in $A_u \neq A_v$ respectively, then using Lemma 2.7:

$$d(b_i, b_j) = d(b_i, e_u) + d(e_u, e_v) + d(e_v, b_j)$$
$$= d(f_u(b_i), f_u(e_u)) + d(e_u, e_v) + d(f_v(e_v), f_v(b_j))$$
$$= d(f(b_i), e_u) + d(e_u, e_v) + d(e_v, f(b_j))$$
$$= d(f(b_i), f(b_j))$$

Therefore the function f is an isometric embedding from $\bigcup\limits_{v=1}^{m} K_v$ to M.

Let $c_j = f(b_j)$ for all $j \in \{1, \ldots, l\}$. Then clearly $d(b_i, b_j) = d(f(b_i), f(b_j)) = d(c_i, c_j)$ for all $i, j \in \{1, \ldots, l\}$. Now let $i \in \{1, \ldots, k\}$ and $j \in \{1, \ldots, l\}$. Let e_v be the closest point to b_i in E_a. Then:

$$d(a_i, b_j) = d(a_i, e_v) + d(e_v, b_j)$$
$$= d(a_i, e_v) + d(f(e_v), f(b_i))$$
$$= d(a_i, e_v) + d(e_v, c_j) = d(a_i, c_j)$$

Thus, the claim is true.

The values of quantifier free formulas in $\mathcal{M}$ are determined by distances in M. (This can be shown using induction on the definition of quantifier free formula, since connectives are continuous functions on atomic formulas, which in this case are all of the form $d(t_1, t_2)$ for terms t_1, t_2.) So, the preceding claim implies that for any quantifier free formula $\varphi(x_1, \ldots, x_k, y_1, \ldots, y_l)$ and any $b_1, \ldots, b_l \in \mathcal{N}$, there exist $c_1, \ldots, c_l \in \mathcal{M}$ so that:

$$\varphi(a_1, \ldots, a_k, b_1, \ldots, b_l)^{\mathcal{N}} = \varphi(a_1, \ldots, a_k, c_1, \ldots, c_l)^{\mathcal{N}} = \varphi(a_1, \ldots, a_k, c_1, \ldots, c_l)^{\mathcal{M}}$$

Then standard arguments about infima imply:

$$\inf_{y_1} \ldots \inf_{y_l} \varphi(a_1, \ldots, a_k, y_1, \ldots, y_l)^{\mathcal{M}} = \inf_{y_1} \ldots \inf_{y_l} \varphi(a_1, \ldots, a_k, y_1, \ldots, y_l)^{\mathcal{N}}$$

Therefore, $\mathcal{M}$ is an existentially closed model of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$. □

8   Properties of the theory $\mathrm{rb}\mathbb{R}\mathrm{T}_r$

In this section we show that $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ has quantifier elimination and is complete and stable, but not superstable. We characterize types, and show that the space of 2-types over the empty set has metric density 2^{ω}. We also characterize definable closure and algebraic closure in models of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$.

8.1 Notation If $A = A_1, \ldots, A_n$ is a finite sequence of subsets of a model, we will use the shorthand notation $A_1 \ldots A_n$ for the union $A_1 \cup \cdots \cup A_n$. If some A_j has a single element, we write the element instead of the set; for example ABC stands for $A \cup B \cup C$ and Ap stands for $A \cup \{p\}$. This notation will be used mainly when we are considering types over sets of parameters.

8.2 Lemma *The* L_r*-theory* rb$\mathbb{R}$T$_r$ *has quantifier elimination.*

Proof By Theorem 7.12, Theorem 6.4 and Proposition 3.8. □

8.3 Corollary *The* L_r*-theory* rb$\mathbb{R}$T$_r$ *is complete.*

Proof In any model of rb$\mathbb{R}$T$_r$ we may embed the structure consisting of just the base point. This fact together with quantifier elimination implies that rb$\mathbb{R}$T$_r$ is complete. □

We now turn our attention to a discussion of types and type spaces. Let $\mathcal{M} \models$ rb$\mathbb{R}$T$_r$. Recall that if $b = b_1, \ldots, b_n$ is tuple of elements in $\mathcal{M}$, then $\mathrm{tp}_{\mathcal{M}}(b/A)$ is the complete n-type of b over A in $\mathcal{M}$. The space of all n-types over $\emptyset$ in models of a theory T is denoted $S_n(T)$. The space of n-types over A is denoted $S_n(T_A)$ or $S_n(A)$ when the theory T is clear from context. If q is an n-type and a is an n-tuple from M such that $\mathcal{M} \models q(a)$ we say *a realizes q in $\mathcal{M}$*, and write $a \models q$. If $\varphi(x_1, \ldots, x_n)$ is an L_r-formula, we write $\varphi(x_1, \ldots, x_n)^q$ for the value of this formula specified by the type q.

The following lemma gives criteria for when two n-tuples have the same type. Recall that for a subset A of a given $\mathbb{R}$-tree M

$$E_A = \bigcup \{[a_1, a_2] \mid a_1, a_2 \in A\}$$

is the smallest subtree of M containing A. The closure $\overline{E_A}$ is the smallest closed subtree containing A.

8.4 Lemma *Let* $\mathcal{M} \models$ rb$\mathbb{R}$T$_r$. *Fix* $A \subseteq M$.

(1) *Let* $b, c \in M$. *Then* $\mathrm{tp}_{\mathcal{M}}(b/A) = \mathrm{tp}_{\mathcal{M}}(c/A)$ *if and only if b and c have the same unique closest point* $e \in \overline{E_{Ap}}$ *and* $d(b, e) = d(c, e)$.
(2) *Let* $b = (b_1, \ldots, b_n)$ *and* $c = (c_1, \ldots, c_n)$ *be tuples in M. Then* $\mathrm{tp}_{\mathcal{M}}(b/A) = \mathrm{tp}_{\mathcal{M}}(c/A)$ *if and only if* $\mathrm{tp}_{\mathcal{M}}(b_i/A) = \mathrm{tp}_{\mathcal{M}}(c_i/A)$ *for all* $1 \leq i \leq n$ *and* $d(b_i, b_j) = d(c_i, c_j)$ *for all* $1 \leq i, j \leq n$.

(3) Let $b = (b_1, \ldots, b_n)$ and $c = (c_1, \ldots, c_n)$ be tuples in M. Let E_b be the subtree
     of M spanned by Apb, and similarly define E_c. Then $\mathrm{tp}_{\mathcal{M}}(b/A) = \mathrm{tp}_{\mathcal{M}}(c/A)$
     if and only if there is an isometry $f \colon E_b \to E_c$ that fixes each element of Ap (or
     equivalently of $\overline{E_{Ap}}$) and satisfies $f(b_i) = c_i$ for all $i = 1, \ldots, n$.

Proof To show (1), first assume $\mathrm{tp}_{\mathcal{M}}(b/A) = \mathrm{tp}_{\mathcal{M}}(c/A)$. By Lemma 2.14(3) this
implies $d(b, a) = d(c, a)$ for all $a \in \overline{E_{Ap}}$, which means b and c must have the same
unique closest point $e \in \overline{E_{Ap}}$. Moreover, $d(b, e)$ must equal $d(c, e)$.

For the other direction, assume b and c have the same unique closest point $e \in \overline{E_{Ap}}$
and that $d(b, e) = d(c, e)$. Since $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ has quantifier elimination, and the values of
quantifier-free formulas are determined by the values of atomic formulas, it suffices to
show $d(a, b) = d(a, c)$ for all $a \in Ap$. This follows easily from our assumptions using
Lemma 2.4, since for any $a \in Ap$, the point e must be on both segments $[a, b]$ and
$[a, c]$.

Statement (2) follows from part (1) and the fact that $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ admits quantifier elim-
ination. Note that $\mathrm{tp}_{\mathcal{M}}(b_i/A) = \mathrm{tp}_{\mathcal{M}}(c_i/A)$ for all $1 \le i \le n$ if and only if
$d(b_i, a) = d(c_i, a)$ for all $a \in Ap$. Again, we use that the values of quantifier-free
formulas are determined by the values of the atomic formulas. The atomic formu-
las involved are all of the form $d(t_1, t_2)$ where t_1 and t_2 are variables $x_1, \ldots, x_n$ or
parameters from Ap.

Statement (3) follows from (2) and Lemma 2.14.   □

With Lemma 8.4 in hand, we can finish the characterization of κ-saturated $\mathbb{R}$-trees,
completing the result promised in Section 7.

8.5 Theorem *Let $\mathcal{M} \models \mathrm{rb}\mathbb{R}\mathrm{T}_r$. Let κ be an infinite cardinal. Then $\mathcal{M}$ is κ-saturated
if and only if M has at least κ-many branches at every point $a \in M$ where $d(p, a) < r$.*

Proof The forward direction is Lemma 7.10. Now, assume M has κ-many branches
at every point $a \in M$ where $d(p, a) < r$. Let $A \subseteq M$ have cardinality $< \kappa$.

Let q be a 1-type over A. By Lemma 8.4, this type is determined by a closest point
$e \in \overline{E_{Ap}}$ and a distance $0 \le s \le r - d(p, e)$, where s is the distance between $\overline{E_{Ap}}$ and
any realization of q. Suppose β is a branch in M at e. The complement of β is path
connected and closed, so if the complement contains all of Ap, then it contains $\overline{E_{Ap}}$.
Hence, if β contains a point in $\overline{E_{Ap}}$, it must contain an element of Ap.

Because there are $\kappa > |Ap|$-many branches at e, there are branches at e in M that do not intersect $\overline{E_{Ap}}$. Let β be one of those branches. Lemma 7.7 allows us to take $x \in \beta$ such that $d(p,x) = r$, and hence $d(x,e) = r - d(p,e)$, since p is not in β. Since $0 \leq s \leq r - d(p,e)$ we may take b on the segment $[e,x]$ with $d(e,b) = s$. Then b satisfies the type q. Therefore, all 1-types over A are realized in $\mathcal{M}$, implying that $\mathcal{M}$ is κ-saturated. $\qquad\square$

Our next result gives a description of the type space $S_n(A)$ with its logic topology, where A is a subset of a model $\mathcal{M}$ of rb$\mathbb{R}$T$_r$. If $q \in S_n(A)$, we take the free variables used in q to be the distinct variables $x_1, \ldots, x_n$ and require that $X = \{x_1, \ldots, x_n\}$ be disjoint from Ap. Note that q induces a pseudometric, which we denote ρ_q, on ApX, by setting $\rho_q(t_1, t_2) = d(t_1, t_2)^q$ for each $t_1, t_2 \in ApX$. That is, if $(b_1, \ldots, b_n)$ realizes $q(x_1, \ldots, x_n)$ in some elementary extension $\mathcal{N}$ of $\mathcal{M}$, and we let $\pi \colon ApX \to N$ be the evaluation map taking each $a \in Ap$ to itself and taking each x_i to b_i, then we are setting $\rho_q(t_1, t_2)$ equal to $d^{\mathcal{N}}(\pi(t_1), \pi(t_2))$. Note that π can be viewed as the quotient map from the pseudometric space (ApX, ρ_q) to the metric space $(Ap\{b_1, \ldots, b_n\}, d^{\mathcal{N}})$, which is 0-hyperbolic since N is an $\mathbb{R}$-tree. It follows that ρ_q is 0-hyperbolic on ApX. (See Definition 2.9 and Lemma 2.14.)

For each $q \in S_n(A)$, we regard ρ_q as an element of the product space $[0, 2r]^{(ApX)^2}$, to which we give the product topology. Let HM(ApX) denote the set of all $\rho \in [0, 2r]^{(ApX)^2}$ such that ρ defines a 0-hyperbolic pseudometric on ApX extending $d^{\mathcal{M}}$ on Ap and satisfying $\rho(p, x_i) \leq r$ for all $i = 1, \ldots, n$. This is a closed subset of $[0, r]^{(ApX)^2}$ and hence it is a compact Hausdorff space in the induced topology.

8.6 Lemma *Let* $\mathcal{M} \models$ rb$\mathbb{R}$T$_r$ *and* $A \subseteq M$.

(1) *The image of the map* $q \mapsto \rho_q$ *defined above on* $S_n(A)$ *is* HM(ApX).

(2) *The map* $q \mapsto \rho_q$ *is a homeomorphism between* $S_n(A)$ *with the logic topology and* HM(ApX) *with the induced topology.*

Proof Let $\mathcal{M} \models$ rb$\mathbb{R}$T$_r$ and let $A \subseteq M$.

(1) Suppose $\rho \in$ HM(ApX). Recall that the quotient metric space of (ApX, ρ) is 0-hyperbolic (see the comment just before Lemma 2.11) and the quotient map from (ApX, ρ) to that metric space is isometric. Therefore, by Lemma 2.14(2) there is an isometric mapping f from (ApX, ρ) into an $\mathbb{R}$-tree (N, d). We may assume that $d(f(p), u) \leq r$ for all $u \in N$ and that (N, d) is complete, and hence that $\mathcal{N} = (N, d, f(p)) \models \mathbb{R}T_r$. By Theorem 7.12 we may assume that $\mathcal{N}$ is a κ-saturated and strongly κ-homogeneous model of rb$\mathbb{R}$T$_r$, where κ is an infinite cardinal $> |A|$. Since

rb$\mathbb{R}$T$_r$ admits QE we may assume that $Ap \subseteq N$ and that f is the identity on Ap. Now let $b = (b_1, \ldots, b_n) = (f(x_1), \ldots, f(x_n)) \in N^n$, and let $q(x_1, \ldots, x_n) = tp(b/A)$ in $\mathcal{N}$. It is easy to see that $\rho = \rho_q$, and hence ρ is in the image of the map $q \mapsto \rho_q$ as desired.

(2) Quantifier elimination for rb$\mathbb{R}$T$_r$ implies that the map $q \mapsto \rho_q$ is 1-1 and continuous for the logic topology on $S_n(A)$ and the induced topology on HM(ApX), and this map is surjective by part (1). Both topologies are compact Hausdorff, and hence the map must be a homeomorphism. $\square$

Next, we give some results about the type spaces as metric spaces. First, a reminder of how the induced metric on types is defined.

Let $\mathcal{M} \models$ rb$\mathbb{R}$T$_r$ be such that every type in $S_n($rb$\mathbb{R}$T$_r))$ is realized in $\mathcal{M}$ for each $n \geq 1$. As defined in Ben Yaacov, Berenstein, Henson and Usvyatsov [4], the d-metric on n-types over the empty set is

$$d(q, q') = \inf_{b \models q,\ b' \models q'} \max_{i=1,\ldots,n} d^{\mathcal{M}}(b_i, b'_i)$$

where $b = b_1, \ldots, b_n$ and $b' = b'_1, \ldots, b'_n$ are tuples in $\mathcal{M}$. Note that since $\mathcal{M}$ realizes all $2n$-types, the infimum in the definition is actually realized. This definition can be extended in the obvious way to spaces of types over parameters.

First we consider the distance between 1-types over A for the theory rb$\mathbb{R}$T$_r$.

8.7 Lemma *Let $\mathcal{M} \models$ rb$\mathbb{R}$T$_r$ and $A \subseteq M$.*

  (1) *As a metric space, $S_1($rb$\mathbb{R}$T$_r) = S_1(\emptyset)$ is isometric to $[0, r]$.*
  (2) *More generally, $S_1(A)$ is in bijective correspondence with the set of ordered pairs*

$$\{(e, s) \mid e \in \overline{E_{Ap}} \quad \text{and} \quad s \in [0, r - d(p, e)] \subseteq \mathbb{R}\}.$$

  Indeed, $q \in S_1(A)$ is associated to (e, s) exactly when e is the unique closest point in $\overline{E_{Ap}}$ to any realization b of q (in an elementary extension $\mathcal{N}$ of $\mathcal{M}$) and $s = d^{\mathcal{N}}(e, b)$; we will denote this distance s by $d(e, x)^q$.
  (3) *Given $q, q' \in S_1(A)$ let $e_q, e_{q'}$ be the unique closest points in $\overline{E_{Ap}}$ to realizations of q, q' respectively.*
    (a) *If $e_q \neq e_{q'}$, then $d(q, q') = d(e_q, x)^q + d(e_q, e_{q'}) + d(e_{q'}, x)^{q'}$.*
    (b) *If $e_q = e_{q'}$, then $d(q, q') = |d(p, x)^q - d(p, x)^{q'}|$.*
  (4) *The metric space $(S_1(A), d)$ is an $\mathbb{R}$-tree, isometric to the $\mathbb{R}$-tree obtained from $\overline{E_{Ap}}$ by amalgamating an additional branch β_e at e that is isometric to the interval $(0, r - d(p, e)]$, for each $e \in \overline{E_{Ap}}$.*

Proof Let $\mathcal{M} \models \mathrm{rb}\mathbb{R}\mathrm{T}_r$ and $A \subseteq M$.

(1) Lemma 8.6 gives a bijection f from $S_1(\emptyset) = S_1(\mathrm{rb}\mathbb{R}\mathrm{T}_r)$ onto $[0, r]$, in which each type $q(x)$ is mapped to $d(p, x)^q$. We show that f is isometric. Given two 1-types $q, q' \in S_1(\mathrm{rb}\mathbb{R}\mathrm{T}_r)$ a configuration minimizing the distance between realizations $b \models q$ and $b' \models q'$ is the one where b, b' and p are arranged along a piecewise segment. That configuration makes $d(b, b') = |d(p, b) - d(p, b')|$ which is the least possible value of $d(b, b')$. by the triangle inequality. Thus, $d(q, q') = d(b, b') = |d(p, b) - d(p, b')| = |f(q) - f(q')|$, showing that f is an isometry.

(2) This is a consequence of Lemmas 8.4, 8.6, and 2.14(3). Lemma 8.4(1) shows that the pair (e, s) associated to q in (2) is indeed determined by q. A description of how to find b realizing a type q associated to (e, s) in a sufficiently saturated elementary extension of $\mathcal{M}$ is given in the last paragraph of the proof of Theorem 8.5.

(3) Let $q, q' \in S_1(A)$. First assume $e_q \neq e_q'$. Then for any realizations $b \models q$ and $b' \models q'$, $[b, e_q, e_{q'}, b']$ is a piecewise segment by Lemma 2.7, and the conclusion follows. If $e_q = e_{q'}$, then as in the proof of (1), a minimizing configuration will have the three points $b \models q$ and $b' \models q'$ and p along a geodesic segment, and the conclusion follows.

(4) This is immediate from (2) and (3). $\qquad\square$

Note that part (2) in Lemma 8.7 gives a description of $S_1(A)$ that is different from the one in Lemma 8.6. It is possible to give an analogous description of $S_n(A)$ when $n > 1$. Let $x_1, \ldots, x_n$ be distinct variables, with $X = \{x_1, \ldots, x_n\}$ disjoint from $\overline{E_{Ap}}$. Suppose we are given $q(x_1, \ldots, x_n) \in S_n(A)$ and a realization $(b_1, \ldots, b_n)$ of q in some model $\mathcal{M}$ of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ containing A. Let π be the function on the set $\overline{E_{Ap}}X$ defined by $\pi(e) = e$ for $e \in \overline{E_{Ap}}$ and $\pi(x_i) = b_i$ for $i = 1, \ldots, n$. According to parts (1) and (2) of Lemma 8.4, the type q is determined by the pseudometric $\eta = \eta_q$ on $\overline{E_{Ap}}X$ that is defined by setting $\eta(t_1, t_2) = d^{\mathcal{M}}(\pi(t_1), \pi(t_2))$ for all $t_1, t_2 \in \overline{E_{Ap}}X$. The key properties of η are the following:

  (a) η agrees with $d^{\mathcal{M}}$ on $\overline{E_{Ap}}$;
  (b) $\eta(p, x_i) \leq r$ for all $i = 1, \ldots, n$;
  (c) $\eta(e, x_i) = \eta(e, e_i) + \eta(e_i, x_i)$ for all $e \in \overline{E_{Ap}}$, $i = 1, \ldots, n$; and
  (d) η is 0-hyperbolic.

For every such pseudometric η there is an n-type $q \in S_n(A)$ for which $\eta = \eta_q$. Indeed, by Lemma 2.14, such a pseudometric η is determined by its restriction to the set ApX, which is an element of $\mathrm{HM}(ApX)$ defined above. Thus it corresponds to a type in $S_n(A)$ by Lemma 8.6.

For $n > 1$ it seems to be very complex to give a precise description, as in the preceding lemma, of the metric on $S_n(A)$ over the theory $\mathrm{rb}\mathbb{R}\mathrm{T}_r$. Accordingly, we limit ourselves to giving (in the next result) a precise statement of the metric density of the space of 2-types over $\emptyset$. Its proof illustrates some of the ideas needed to understand these metric spaces more completely.

8.8 Proposition *The space of 2-types over the empty set has metric density character equal to 2^ω.*

Proof It is clear that $S_2(\mathrm{rb}\mathbb{R}\mathrm{T}_r)$ has cardinality 2^ω. Hence it suffices to find, for some $\delta > 0$, a set of 2^ω-many types in $S_2(\mathrm{rb}\mathbb{R}\mathrm{T}_r)$ that are pairwise at distance $\geq \delta$.

For each $s \in [0, r]$ let $q_s(x, y) \in S_2(\mathrm{rb}\mathbb{R}\mathrm{T}_r)$ contain the conditions $d(p, x) = d(p, y) = r$ and $d(x, y) = 2s$. Note that for every s there is a unique such type. If (a_s, b_s) realizes q_s in $\mathcal{M} \models \mathrm{rb}\mathbb{R}\mathrm{T}_r$ we let $Y_s = Y(p, a_s, b_s)$, and note that $d(a_s, Y_s) = d(b_s, Y_s) = s$ and $d(p, Y_s) = r - s$.

Suppose $0 \leq t < s \leq r$. We will show that the distance between q_s and q_t is $\geq 2t$. So if we take $0 < \delta < r$, the types in $\{q_s \mid \delta \leq s \leq r\}$ are pairwise at distance $\geq 2\delta$, giving the desired result.

So suppose for $0 \leq t < s \leq r$ we have (a_s, b_s) realizing q_s and (a_t, b_t) realizing q_t, in a model $\mathcal{M}$ of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$, and that $d(a_s, a_t) \leq 2t$. Letting $c = Y(p, a_s, a_t)$ we have $d(p, c) \geq r - t$, implying that Y_s and Y_t are both in the geodesic segment $[p, c]$. Then, since $Y_s \neq Y_t$, we have by Lemma 2.7 that $[b_s, Y_s, Y_t, b_t]$ is a piecewise segment, and hence

$$d(b_s, b_t) = d(b_s, Y_s) + d(Y_s, Y_y) + d(Y_t, b_s) = s + d(Y_s, Y_t) + t > 2t$$

which completes the proof since the realizations of q_s, q_t were arbitrary.   □

Remark A similar argument shows that if X is a metric $\mathbb{R}$-tree that contains an isometric copy of every 4-point 0-hyperbolic metric space of diameter at most r, then the density character of X is at least 2^ω. Likewise, if a is a point different from p in $\mathcal{M} \models \mathrm{rb}\mathbb{R}\mathrm{T}_r$, then the type space $S_1(\{a\})$ (relative to $\mathrm{rb}\mathbb{R}\mathrm{T}_r$) has density character exactly 2^ω.

Next, we show that the definable closure and the algebraic closure of a set of parameters A are the same, and are equal to the closed subtree spanned by Ap.

8.9 Proposition *Let $\mathcal{M} \models \mathrm{rb}\mathbb{R}\mathrm{T}_r$ with (M, d, p) as its underlying $\mathbb{R}$-tree. Let $A \subseteq M$ be a set of parameters. Then $\mathrm{dcl}(A) = \mathrm{acl}(A) = \overline{E_{Ap}}$.*

Proof If $a, b \in Ap$, then by Lemma 5.2 any point in $[a, b]$ is in $\mathrm{dcl}(A)$. Then $E_{Ap} \subseteq \mathrm{dcl}(A)$ because E_{Ap} is the union of all such geodesic segments. Since $\mathrm{dcl}(A)$ is closed, $\overline{E_{Ap}} \subseteq \mathrm{dcl}(A)$. Combined with the fact that $\mathrm{dcl}(A) \subseteq \mathrm{acl}(A)$, this gives $\overline{E_{Ap}} \subseteq \mathrm{dcl}(A) \subseteq \mathrm{acl}(A)$. It remains to show $\mathrm{acl}(A) \subseteq \overline{E_{Ap}}$, which we do in the contrapositive.

If $c \notin \overline{E_{Ap}}$, let e be the unique closest point to c in $\overline{E_{Ap}}$. This e exists by Lemma 2.5. Let β be the branch at e that contains c. Using Lemma 6.1, construct an extension $\mathcal{N} \models \mathrm{rb}\mathbb{R}\mathrm{T}_r$ of $\mathcal{M}$ which adds an infinite number of branches at e, each of which is isometric to β. By Lemma 8.4, on each of these branches is a realization of $\mathrm{tp}(c/A)$. The distance between any two such realizations is $2d(e, c) > 0$. This gives us a non-compact set of realizations of $\mathrm{tp}(c/A)$. Therefore, $c \notin \mathrm{acl}(A)$. $\qquad\square$

To finish this section we show $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ is stable, but not superstable (ie, it is strictly stable).

8.10 Theorem *The theory* $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ *is stable. Indeed when κ is an infinite cardinal,* $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ *is κ-stable if and only if κ satisfies $\kappa^\omega = \kappa$.*

Proof Let κ be an infinite cardinal. Let $\mathcal{M} \models \mathrm{rb}\mathbb{R}\mathrm{T}_r$ be κ^+-saturated, with underlying $\mathbb{R}$-tree (M, d, p).
First, assume $\kappa = \kappa^\omega$. Let $|A| = \kappa$. Then:

$$|E_{Ap}| \leq |A \times A| 2^\omega = \kappa^2 2^\omega \leq \kappa^\omega 2^\omega = \kappa^\omega = \kappa$$

Thus, $|E_{Ap}| = \kappa$. We count possible 1-types using Lemma 8.7, showing that:

$$|S_1(A)| \leq |\overline{E_{Ap}} \times [0, r]| = |\overline{E_{Ap}}| 2^\omega \leq |E_{Ap}|^\omega 2^\omega = \kappa^\omega 2^\omega = \kappa^\omega = \kappa$$

Thus $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ is κ-stable.

For the other direction, assume $\kappa < \kappa^\omega$. We construct, via a tree construction, a subset A of M with $|A| = \kappa$ and $|\overline{E_{Ap}}| = \kappa^\omega$. At Step 1, we choose κ-many points $(a_i \mid i < \kappa)$ on distinct branches at p, each with distance $r/4$ from p. We can do this since there are at least κ branches of sufficient height at every point in M by Theorem 8.5 and Lemma 7.7. Note that we only need κ-saturation to guarantee κ-many branches. At Step 2, for each a_i we choose κ-many points on distinct branches at a_i, each with distance $r/8$ from a_i, and distance $3r/8$ from p. We index these points by $(a_{i,j} \mid i, j < \kappa)$. At Step n for $n \geq 2$, we have already designated points $a_{i_1, i_2, \ldots, i_{n-1}}$ each of which has distance $\sum_{k=1}^{n-1} r/2^{k+1}$ from p. At each of these points choose κ-many points on distinct branches at $a_{i_1, i_2, \ldots, i_{n-1}}$, each with distance $r/2^{n+1}$

from $a_{i_1,i_2,\dots,i_{n-1}}$, and distance $\sum_{k=1}^{n} r/2^{k+1}$ from p. We index these new points by $(a_{i_1,i_2,\dots,i_n} \mid i_1, \dots, i_n < \kappa)$. Let $A = \bigcup_{k=1}^{\infty}(a_{i_1,\dots,i_k} \mid i_1, \dots, i_k < \kappa)$. If we associate p with the empty sequence, then the elements of Ap are in 1-1 correspondence with $\kappa^{<\omega}$. So, the cardinality of A is $|\kappa^{<\omega}| = \kappa$.

If σ and τ are non-empty finite sequences from κ, we have that $[p, a_\sigma, a_{\sigma,\tau}]$ is a piecewise segment in M. (Here σ, τ denotes the concatenation of σ and τ.) Moreover, for every such σ and every distinct $i, j < \kappa$, the points $a_{\sigma,i}$, $a_{\sigma,j}$, and p are in distinct branches at a_σ.

For each function $f: \omega \to \kappa$ with $f(0) = 0$ we define a sequence (b_n^f) of elements of A by setting $b_0^f = p$ and $b_n^f = a_{f(1),\dots,f(n)}$ for $n > 0$. For each $n \in \omega$ we have that $[b_0^f, \dots, b_n^f]$ is a piecewise segment in M. Further, b_k^f has distance $r/2^{k+1}$ from b_{k-1}^f for every $k > 1$, making (b_n^f) a Cauchy sequence. Since M is complete, b_n^f must converge to a limit, which we denote by u_f. Note that $d(p, u_f) = r/2$; indeed, $d(b_n^f, u_f) = r/2^{n+1}$ for every n.

Let f and g be two such functions from ω to κ, and suppose $m > 0$ is the first index at which $f(m)$ and $g(m)$ disagree. An easy calculation shows that u_f, u_g, and p are in distinct branches at $b_{m-1}^f = b_{m-1}^g$ in M and $d(u_f, u_g) = r/2^{m-1}$. In particular, the map $f \mapsto u_f$ is injective, so the set of all points u_f has cardinality κ^ω.

For each u_f as above, let β_f be the branch at u_f in M that contains p. Evidently $A \subseteq \beta_f$; since $\beta_f \cup \{u_f\}$ is closed and path-connected, it must contain $\overline{E_{Ap}}$. We choose a point $v_f \in M$ in a different branch at u_f from β_f such that $d(v_f, u_f) = r/2$. Note that u_f is the closest point to v_f in $\overline{E_{Ap}}$.

By Lemma 8.4 the type of v_f over A is determined by u_f and $d(v_f, u_f)$. Further, by Lemma 8.7(3), for distinct f, g we have:

$$d(\text{tp}(v_f/A), \text{tp}(v_g/A)) = d(u_f, v_f) + d(u_f, u_g) + d(u_g, v_g) > r$$

Therefore, the metric density character of $S_1(A)$ is at least κ^ω. Since A has cardinality κ and $\kappa^\omega > \kappa$, this shows that the theory $\text{rb}\mathbb{R}\text{T}_r$ is not κ-stable. $\square$

9   The independence relation for $\text{rb}\mathbb{R}\text{T}_r$

In this section we characterize the model theoretic independence relation of $\text{rb}\mathbb{R}\text{T}_r$ and show that in models of $\text{rb}\mathbb{R}\text{T}_r$, types have canonical bases that are easily-described sets of ordinary (not imaginary) elements.

Let κ be a cardinal so that $\kappa = \kappa^\omega$ and $\kappa > 2^\omega$. In this section let $\mathbb{U}$ be a κ-universal domain for rb$\mathbb{R}$T$_r$. (That is, $\mathbb{U}$ is κ-saturated and κ-strongly homogeneous; see Ben Yaacov, Berenstein, Henson and Usvyatsov [4, Definition 7.13] and the discussion preceding it.) A subset of $\mathbb{U}$ is *small* if its cardinality is $< \kappa$.

9.1 Definition Let A, B and C be small subsets of $\mathbb{U}$. Say A is *-independent from B over C, denoted $A \underset{C}{\overset{*}{\smile}} B$, if and only if for all $a \in A$ we have $\mathrm{dist}(a, \overline{E_{BCp}}) = \mathrm{dist}(a, \overline{E_{Cp}})$.

9.2 Lemma $A \underset{C}{\overset{*}{\smile}} B$ if and only if for all $a \in A$ the closest point to a in $\overline{E_{BCp}}$ is the same as the closest point to a in $\overline{E_{Cp}}$.

Proof ($\Rightarrow$) Assume $A \underset{C}{\overset{*}{\smile}} B$. Take an arbitrary $a \in A$. Let e_1 be the unique closest point to a in $\overline{E_{BCp}}$ and e_2 the unique closest point to a in $\overline{E_{Cp}}$. We assumed $\mathrm{dist}(a, \overline{E_{BCp}}) = \mathrm{dist}(a, \overline{E_{Cp}})$, which implies $d(a, e_1) = d(a, e_2)$. Since $e_2 \in \overline{E_{Cp}} \subseteq \overline{E_{BCp}}$, we know $e_1 \in [a, e_2]$ by Lemma 2.5. Therefore, $e_1 = e_2$. Since a was arbitrary, we know this holds for all $a \in A$.

($\Leftarrow$) Assume for all $a \in A$ the closest point to a in $\overline{E_{BCp}}$ is the closest point to a in $\overline{E_{Cp}}$. Then clearly $\mathrm{dist}(a, \overline{E_{BCp}}) = \mathrm{dist}(a, \overline{E_{Cp}})$ for all $a \in A$. $\qquad\square$

9.3 Theorem *The relation $\overset{*}{\smile}$ is the model theoretic independence relation for* rb$\mathbb{R}$T$_r$. *Moreover, types over arbitrary sets of parameters are stationary.*

Proof We will show $\overset{*}{\smile}$ satisfies all the properties of a stable independence relation on a universal domain of a stable theory as given in Ben Yaacov, Berenstein, Henson and Usvyatsov [4, Theorem 14.12]. Then by [4, Theorem 14.14] we know $\overset{*}{\smile}$ is the model theoretic independence relation for the stable theory rb$\mathbb{R}$T$_r$.

(1) Invariance under automorphisms

Any automorphism σ of $\mathbb{U}$ satisfies $\sigma(\overline{E_{Ap}}) = \overline{E_{\sigma(Ap)}}$ and is distance preserving.

(2) Symmetry: if $A \underset{C}{\overset{*}{\smile}} B$ then $B \underset{C}{\overset{*}{\smile}} A$.

Assume $A \underset{C}{\overset{*}{\smile}} B$. This means for all $a \in A$ we have that the closest point in $\overline{E_{BCp}}$ to a is $e_a \in \overline{E_{Cp}}$. Thus, by Lemma 2.4, for any $a \in A$, for any $y \in \overline{E_{BCp}}$ we have $[a, y] \cap \overline{E_{Cp}} \neq \emptyset$. It follows that for any $x \in \overline{E_{Ap}}$, for any $y \in \overline{E_{BCp}}$ there exists a point of $\overline{E_{Cp}}$ on $[x, y]$. Let $b \in B$. Then for any $x \in \overline{E_{Ap}}$ there is a point of $\overline{E_pC}$ on $[x, b]$. It follows that the closest point in $\overline{E_{ACp}}$ to any $b \in B$ is in $\overline{E_{Cp}}$.

(3) Transitivity: $A \underset{C}{\overset{*}{\downarrow}} BD$ if and only if $A \underset{C}{\overset{*}{\downarrow}} B$ and $A \underset{BC}{\overset{*}{\downarrow}} D$.

We know $E_{Cp} \subseteq E_{BCp} \subseteq E_{BCDp}$ which implies:

$$\mathrm{dist}(a, \overline{E_{Cp}}) \geq \mathrm{dist}(a, \overline{E_{BCp}}) \geq \mathrm{dist}(a, \overline{E_{BCDp}})$$

Therefore $\mathrm{dist}(a, \overline{E_{BCDp}}) = \mathrm{dist}(a, \overline{E_{Cp}})$ if and only if

$$\mathrm{dist}(a, \overline{E_{BCp}}) = \mathrm{dist}(a, \overline{E_{Cp}}) \quad \text{and} \quad \mathrm{dist}(a, \overline{E_{BCDp}}) = \mathrm{dist}(a, \overline{E_{BCp}}).$$

Hence

$$A \underset{C}{\overset{*}{\downarrow}} BD \quad \text{if and only if} \quad A \underset{C}{\overset{*}{\downarrow}} B \quad \text{and} \quad A \underset{BC}{\overset{*}{\downarrow}} D.$$

(4) Finite character: $A \underset{C}{\overset{*}{\downarrow}} B$ if and only if $a \underset{C}{\overset{*}{\downarrow}} B$ for all finite tuples $a \in A$.

This is clear from the definition.

(5) Extension: for all A, B, C we can find A' such that $\mathrm{tp}(A/C) = \mathrm{tp}(A'/C)$ and $A' \underset{C}{\overset{*}{\downarrow}} B$.

By finite character and saturation of $\mathbb{U}$, it suffices to show this statement when A is a finite tuple. Let $e \in \overline{E_{Cp}}$ be the unique point closest to $\overline{E_{Ap}} = E_{Ap}$. Let $\tau < \kappa$ be the cardinality of $\overline{E_{Bp}}$. Then there are at most τ-many branches in $\overline{E_{Bp}}$ at any point of $\overline{E_{Bp}}$. Since A is finite, use Lemma 7.11 to embed a copy of $\overline{E_{Ap}} = E_{Ap}$ on branches at e that do not intersect $\overline{E_{Bp}}$. The image of A under this embedding gives us A'.

(6) Local Character: if $a = a_1, \ldots, a_m$ is a finite tuple, there is a countable $B_0 \subseteq B$ such that $a \underset{B_0}{\overset{*}{\downarrow}} B$.

Let e_i be the closest point of $\overline{E_p B}$ to a_i for $i = 1, \ldots, m$. Let B_i be a countable subset of B such that e_i is an element of $\overline{E_{B_i p}}$. Let $B_0 = \bigcup_i^m B_i$.

(7) Stationarity (over arbitrary sets of parameters): if $\mathrm{tp}(A/C) = \mathrm{tp}(A'/C)$, $A \underset{C}{\overset{*}{\downarrow}} B$, and $A' \underset{C}{\overset{*}{\downarrow}} B$, then $\mathrm{tp}(A/BC) = \mathrm{tp}(A'/BC)$, where C is a small submodel of $\mathbb{U}$.

By quantifier elimination, $\mathrm{tp}(A/BC)$ is determined by $\{\mathrm{tp}(a/BC) \mid a \in A\}$ plus the information $\{d(a_1, a_2) \mid a_1, a_2 \in A\}$. These distances $\{d(a_1, a_2) \mid a_1, a_2 \in A\}$ are fixed by $\mathrm{tp}(A/C)$. Thus, it suffices to show the conclusion in the case when $A = \{a\}$ and $A' = \{a'\}$. If a or a' is in Cp the conclusion is obvious, so assume $a, a' \notin Cp$. The type of a (or a') over BC is determined by two parameters, the unique point in $\overline{E_{BCp}}$ that is closest to a, and the distance from a to that point. Since $a \underset{C}{\overset{*}{\downarrow}} B$, it follows that the closest point in $\overline{E_{Cp}}$ to a is the same as the closest point in $\overline{E_{BCp}}$ to a, and the same is true for a'. Since $\mathrm{tp}(a/C) = \mathrm{tp}(a'/C)$, we know a and a' have the same closest point e in $\overline{E_{Cp}}$ and $d(a, e) = d(a', e)$. Since e is also the closest point in $\overline{E_{BCp}}$ to a and a', we know that $\mathrm{tp}(a/BC) = \mathrm{tp}(a'/BC)$ by Lemma 8.4. $\square$

Canonical Bases

A canonical base of a stationary type is a minimal set of parameters over which that type is definable. However, to avoid a discussion of definable types, we here use an equivalent definition of canonical base, as given in Ben Yaacov, Berenstein, and Henson [3]. As in that paper, we here take advantage of the fact (Lemma 8.9 and part (7) of the proof of Theorem 9.3) that every type over an arbitrary set of parameters is stationary.

For stable theories in general, canonical bases exist as sets of imaginary elements, however, in models of rb$\mathbb{R}$T$_r$, they are sets of ordinary elements. That is, the theory has built-in canonical bases. Indeed, in this setting they are very simple.

For sets $A \subseteq B \subseteq \mathbb{U}$, and $q \in S_n(A)$ we say $q' \in S_n(B)$ is a *non-forking extension* of q if $b \models q'$ implies $b \models q$ and $b \underset{A}{\bigcup} B$. By the definition of independence, the condition $b \underset{A}{\bigcup} B$ implies that the points $e_1, \ldots, e_n$ in $\overline{E_{Ap}}$ closest to $b_1, \ldots, b_n$ respectively must also be the closest points to $b_1, \ldots, b_n$ in $\overline{E_{Bp}}$. Because rb$\mathbb{R}$T$_r$ is stable and all types are stationary, non-forking extensions are unique. Denote the unique non-forking extension of q to the set B by $q{\restriction}^B$. Given a type q over a set $A \subseteq \mathbb{U}$ and an automorphism f of $\mathbb{U}$, $f(q)$ denotes the set of L_r-conditions over $f(A)$ corresponding to the conditions in q, where each parameter $a \in A$ is replaced by its image $f(a)$.

The following is Definition 6.1 from Ben Yaacov, Berenstein, and Henson [3].

9.4 Definition A *canonical base* $Cb(q/A)$ for a type $q \in S_n(A)$ is a subset $C \subseteq \mathbb{U}$ such that for every automorphism $f \in \text{Aut}(\mathbb{U})$, we have: $q{\restriction}^{\mathbb{U}} = f(q){\restriction}^{\mathbb{U}}$ if and only if f fixes each member of C.

The following result describes canonical bases in rb$\mathbb{R}$T$_r$.

9.5 Theorem Let $b = (b_1, \ldots, b_n) \in \mathbb{U}^n$ and $A \subseteq \mathbb{U}$ a set of parameters. Let $q \in S_n(A)$ be the type over A of the tuple b. Then a canonical base of q is given by the set $\{e_i \mid 1 \leq i \leq n\}$, where $e_i \in \overline{E_{Ap}}$ is the closest point to b_i in $\overline{E_{Ap}}$. Note that this set depends only on q.

Proof Let b, $A \subseteq \mathbb{U}$ and $q \in S_n(A)$ be as described in the statement of the theorem. Let $C = \{e_i \mid 1 \leq i \leq n\}$ where $e_i \in \overline{E_{Ap}}$ is the closest point to b_i in $\overline{E_{Ap}}$. First, assume f is an automorphism of $\mathbb{U}$ fixing C pointwise. Let $c = (c_1, \ldots, c_n)$ be a realization of $q{\restriction}^{\mathbb{U}}$ (in some extension of $\mathbb{U}$). Then $c \models q$ and $c \underset{A}{\bigcup} \mathbb{U}$. To show $f(q){\restriction}^{\mathbb{U}} = q{\restriction}^{\mathbb{U}}$ it suffices to show that $c \models f(q)$ and $c \underset{f(A)}{\bigcup} \mathbb{U}$, because then $q{\restriction}^{\mathbb{U}}$ is the

unique non-forking extension of $f(q)$ to $\mathbb{U}$. By Lemma 8.4, an n-type over a set A is determined by the values it assigns to the formulas $d(x_i, x_j)$ and $d(x_i, a)$ for $a \in Ap$. Note that in $f(q)$, the parameter-free L_r-conditions are the same as in the type q. So, for example, $d(x_i, x_j)$ must have the same value in $f(q)$ as in q. Thus, to show that $c \models f(q)$ we just need to show that $d(c_i, a) = d(c_i, f(a))$ for all $a \in Ap$ and $i \in \{1, \ldots, n\}$.

We know $c \models q$. Thus Lemma 8.4 implies that e_i must be the closest point to c_i in $\overline{E_{Ap}}$. Also, $c \underset{A}{\bigcup} \mathbb{U}$ implies that e_i is also the closest point in $\mathbb{U}$ to c_i. Since $f(\overline{E_{Ap}}) \subseteq \mathbb{U}$ we know the closest point in $f(\overline{E_{Ap}})$ to c_i is $e_i = f(e_i)$. Therefore by Lemmas 2.4 and 2.5, we know $d(c_i, a) = d(c_i, e_i) + d(e_i, a)$ and $d(c_i, f(a)) = d(c_i, e_i) + d(e_i, f(a))$ for any $a \in Ap$. Thus,

$$
\begin{aligned}
d(c_i, a) &= d(c_i, e_i) + d(e_i, a) \\
&= d(c_i, e_i) + d(f(e_i), f(a)) \\
&= d(c_i, e_i) + d(e_i, f(a)) = d(c_i, f(a))
\end{aligned}
$$

establishing that $c \models f(q)$.

Since f is an isometry, clearly $f(e_i) = e_i$ is the closest point to c_i in $f(\overline{E_{Ap}})$. The closed subtree $f(\overline{E_{Ap}})$ is equal to the closed subtree $\overline{E_{f(A)p}}$ since $f([a, b]) = [f(a), f(b)]$ for all $a, b \in \mathbb{U}$ and $f(p) = p$. This implies that $c \underset{f(A)}{\bigcup} \mathbb{U}$. We conclude that $f(q){\restriction}^{\mathbb{U}} = q{\restriction}^{\mathbb{U}}$.

For the other direction, assume f is an automorphism of $\mathbb{U}$ that does not fix all of the elements of C. Without loss of generality, assume $f(e_1) \neq e_1$. Let $(c_1, \ldots, c_n) \models q{\restriction}^{\mathbb{U}}$. Then the closest point in $\mathbb{U}$ to c_1 is e_1, which is also the closest point to c_1 in $\overline{E_{Ap}}$. Then, since $f(e_1) \in \mathbb{U}$, the point e_1 must be on the geodesic segment joining $f(e_1)$ and c_1, so $d(f(e_1), c_1) = d(f(e_1), e_1) + d(e_1, c_1)$. Thus, $d(f(e_1), e_1) = d(f(e_1), c_1) - d(e_1, c_1)$. Since $d(f(e_1), e_1) \neq 0$, then $d(f(e_1), c_1) \neq d(e_1, c_1)$. But, $d(e_1, c_1)$ is the value of the formula $d(e_1, x_1)$ in $q{\restriction}^{\mathbb{U}}$, and by the definition of $f(q)$, the value of $d(f(e_1), x_1)$ in $f(q)$ must equal the value of $d(e_1, x_1)$ in q. So, the L-condition $|d(f(e_1), x_1) - d(e_1, c_1)| = 0$ is in the type $f(q)$, and therefore in the type $f(q){\restriction}^{\mathbb{U}}$. Thus, the tuple $c = (c_1, \ldots, c_n)$ cannot be a realization of $f(q){\restriction}^{\mathbb{U}}$, and therefore $q{\restriction}^{\mathbb{U}} \neq f(q){\restriction}^{\mathbb{U}}$. $\square$

10   Models of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$: Examples

In this section we discuss examples of models of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ from the literature. Our first examples come from the explicitly described universal $\mathbb{R}$-trees that are treated in Dyubina and Polterovich [13]. We show that they give exactly the (fully) saturated

models of rb$\mathbb{R}$T$_r$. Our second examples come from asymptotic cones of hyperbolic finitely generated groups. They give exactly the unique saturated model of rb$\mathbb{R}$T$_r$ of density 2^ω.

We begin with a lemma about the density of a κ-saturated model.

10.1 Lemma *Let $\mathcal{M} \models$ rb$\mathbb{R}$T$_r$ with underlying $\mathbb{R}$-tree (M, d, p), and let κ be an infinite cardinal.*

(1) *If there exists $a \in M$ with degree κ, then the density character of M is at least κ.*

(2) *If $\mathcal{M}$ is κ-saturated, then the density character of $\mathcal{M}$ is at least κ^ω.*

Proof (1) If $d(p, a) = r$, then there is a single branch at a, namely the branch containing p. So, we must have $d(p, a) < r$. Using Lemma 7.7 and taking points on different branches at a each with distance r from p, we find a collection of κ-many points such that the distance between any two of them is $2(r - d(p, a))$. Thus, the density character of M must be at least κ.

(2) By Theorem 8.5, at each point in the underlying $\mathbb{R}$-tree of M of $\mathcal{M}$ there are at least κ-many branches. The tree construction from Theorem 8.10 then yields at least κ^ω-many distinct points with pairwise distances at least r. $\qquad\qquad\square$

Universal $\mathbb{R}$-trees

Next, we review the description of the (isometrically) universal $\mathbb{R}$-trees from Dyubina and Polterovich [13] and relate them to saturated models of rb$\mathbb{R}$T$_r$.

10.2 Definition Let μ be a cardinal. An $\mathbb{R}$-tree M is called μ-*universal* if, for any $\mathbb{R}$-tree N with $\leq \mu$ branches at every point, there is an isometric embedding of N into M.

10.3 Example ([13, Example 1.1.1 and Lemma 2.1.1]) For each $\mu \geq 2$ let C_μ be a set with cardinality μ if μ is infinite, and cardinality $\mu - 1$ if μ is finite. Let A_μ be the set of functions $f : (-\infty, \rho_f) \to C_\mu$ from an arbitrary left-infinite open interval to C_μ, where f satisfies:

(1) There exists $\tau_f \leq \rho_f$ so that $f = 0$ for all $t \in (-\infty, \tau_f)$.

(2) The function f is piecewise constant from the right. That is, for any $t \in (-\infty, \rho_f)$ there exists $\delta > 0$ so that f is constant on $[t, t + \delta]$.

On A_μ define the metric d for distinct f, g by $d(f, g) = (\rho_f - s) + (\rho_g - s)$ where $s = \sup\{t \mid f(t') = g(t') \,\forall t' < t\}$.

Dyubina and Polterovich [13] show that A_μ is a complete $\mathbb{R}$-tree with μ-many branches of infinite height at every point. They also show that A_μ is the unique (up to isometry) $\mathbb{R}$-tree with μ branches at each point and that it is homogeneous and μ-universal. In our terminology, for $\mu \geq 3$ the space A_μ is a complete, unbounded richly branching $\mathbb{R}$-tree.

Let $\mathcal{M}_\mu$ be the pointed $\mathbb{R}$-tree obtained by choosing an arbitrary point $p \in A_\mu$ to be the base point and taking $\mathcal{M}_\mu$ to be the closed ball in A_μ of radius r centered at p. We see that $\mathcal{M}_\mu$ is a model of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$. When μ is infinite, $\mathcal{M}_\mu$ is μ-saturated by Theorem 8.5 and $\mathcal{M}_\mu$ has density at least μ^ω by Lemma 10.1(2).

Classical model theory suggests that since $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ is complete and κ-stable exactly when $\kappa = \kappa^\omega$, there should be a κ-saturated model of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ with density κ exactly when $\kappa = \kappa^\omega$. Lemma 10.1(2) shows that if $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ has a κ-saturated model of density κ, then we must have $\kappa = \kappa^\omega$. We verify here that when $\kappa = \kappa^\omega$ and we apply the construction described in the preceding paragraph to $\mu = \kappa$, we do get the unique κ-saturated model $\mathcal{M}_\kappa$ of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ of density κ.

To complete the verification, we need to show that when μ is infinite, the density of $\mathcal{M}_\mu$ is at most μ^ω. In fact, an examination of the construction shows that the cardinality of A_μ is at most μ^ω. Conditions (1) and (2) in Example 10.3 imply that each $f \in A_\mu$ can only change values at a countable or finite number of points. Indeed, given $f \in A_\mu$, we can inductively build a sequence $\{t_n\}_{n\in\mathbb{N}}$ recording where f changes value as follows. For $f \in A_\mu$, let $t_0 = \tau_f$.

Given t_n, let $t_{n+1} = \sup\{t \in (-\infty, \rho_f) \mid f(t) = f(t_n)\}$ if this supremum is $< \rho_f$. If $\sup\{t \in (-\infty, \rho_f) \mid f(t) = f(t_n)\} = \rho_f$, then set $t_{n+i} = \rho_f$ for all $i \in \mathbb{N}$.

- If $t_n = t_{n+1}$, then $t_{n+i} = t_n = \rho_f$ for all $i \in \mathbb{N}$.
- f is constant on $[t_n, t_{n+1})$ for each $n \in \mathbb{N}$.

For each $n \in \mathbb{N}$, let $\alpha_n \in C_\mu$ be the value of f on $[t_n, t_{n+1})$. This gives us a sequence $\{\alpha_n\}$ of elements of C_μ. The sequences $\{t_n\}$ and $\{\alpha_n\}$ determine f. Thus, the cardinality of A_μ is at most $2^\omega \cdot \mu^\omega = \mu^\omega$.

It follows that the metric density of $\mathcal{M}_\mu$ is exactly μ^ω. Thus, in the case that $\mu = \mu^\omega$, the model $\mathcal{M}_\mu$ is a μ-saturated model of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ of density μ, and it is the unique model with those properties.

In the case that $\mu = \mu^\omega$, we can get an alternative model-theoretic argument that A_μ is the unique complete $\mathbb{R}$-tree with μ branches at every point. Given any two complete

$\mathbb{R}$-trees M_1 and M_2 with μ-many branches every point, select a base point in each. Then for each $r > 0$ the closed r-balls in M_1 and M_2 (centered at their respective base points) are saturated models of rb$\mathbb{R}$T$_r$. Hence those r-balls are isomorphic by the fact that saturated models of a complete theory with the same density are isomorphic. A back-and-forth argument can be used to build an isomorphism from M to N, where each time we extend the partial isomorphism to a new point we take its distance from the base point into account and work in closed r-balls for large enough r.

Asymptotic Cones

A finitely generated group is *hyperbolic* if its Cayley graph is a δ-hyperbolic metric space for some $\delta > 0$. A *non-elementary* hyperbolic group is one that has no cyclic subgroup of finite index.

10.4 Definition Let (M, d, p) be a metric space. Let U be a non-principal ultrafilter on $\mathbb{N}$ and let $(\nu_m)_{m \in \mathbb{N}}$ be a sequence of positive integers such that $\lim_{m \to \infty} \nu_m = \infty$. The asymptotic cone of (M, d, p) with respect to $(\nu_m)_{m \in \mathbb{N}}$ and U is the ultraproduct of pointed metric spaces $\prod_U (M, d/\nu_m, p)$. Denote this asymptotic cone by $\mathrm{Con}_{U,(\nu_m)}(M, d, p)$. Elements of $\mathrm{Con}_{U,(\nu_m)}(M, d, p)$ are denoted $[a_n]$ where $a_n \in M$ for each n.

There are versions of this definition that allow, for example, a different choice of base point in each factor. Keeping the same base point is sufficient for our discussion.

10.5 Example An asymptotic cone $\mathrm{Con}_{U,(\nu_m)}(G)$ of a finitely generated group is defined to be the asymptotic cone of its Cayley graph with base point e and some designated word metric on G. It is a fact that any asymptotic cone of a hyperbolic group is an $\mathbb{R}$-tree and is homogeneous (see van den Dries and Wilkie [9] or Drutu [10]). In fact, in the case of a non-elementary hyperbolic group, all asymptotic cones are homogeneous with 2^ω branches at every point (see [10, Proposition 3.A.7]) and are thus are isometric to A_{2^ω} from Example 10.3. A proof of this fact is given below (Lemma 10.7).

10.6 Fact Say B and C are both finite generating sets for the hyperbolic group G and U is a non-principal ultrafilter on $\mathbb{N}$. The word metrics d_B and d_C are Lipschitz equivalent (and the corresponding Cayley graphs are quasi-isometric). It follows that the asymptotic cones $\mathrm{Con}_{U,(\nu_m)}(G, d_B, e)$ and $\mathrm{Con}_{U,(\nu_m)}(G, d_C, e)$ are homeomorphic.

10.7 Lemma *Let G be a non-elementary hyperbolic group. Let U be a non-principal ultrafilter on $\mathbb{N}$ and let $\{\nu_m\}_{m \in \mathbb{N}}$ be a sequence of positive integers such that $\lim_{m \to \infty} \nu_m = \infty$. Then for any finite generating set C for G, the asymptotic cone $\mathrm{Con}_{U,(\nu_m)}(G, d_C, e)$ is a homogeneous richly branching $\mathbb{R}$-tree with 2^ω-many branches at every point.*

Proof Since G is finitely generated, we know that G is countable. Therefore, any asymptotic cone of G has cardinality (and therefore density) at most 2^ω. By Lemma 10.1 we conclude that any point in the cone can have at most 2^ω branches. We also know $\mathrm{Con}_{U,(\nu_m)}(G, d_C, e)$ is a homogeneous $\mathbb{R}$-tree. Since it is non-elementary G contains a free subgroup F with 2 generators (See Bridson and Haefliger [7]). By Fact 10.6 we may assume without loss of generality that the generators of F are also generators of G and that C is a minimal set of generators. In F, for each $m \in \mathbb{N}$, we can find a finite set C_m such that $d(e, a) \geq \nu_m$ for any $a \in C_m$ and so that $(a \cdot b)_e \leq \sqrt{\nu_m}$ for distinct $a, b \in C_m$. The Cayley graph of F is a subgraph of the Cayley graph of G, and the Cayley graph of a free group is an $\mathbb{R}$-tree, thus we use $\mathbb{R}$-tree terminology to describe how to find C_m.

Given $m \in \mathbb{N}^{>0}$ let n_m be the largest integer so that $n_m \leq \sqrt{\nu_m}$. Note that $n_m \to \infty$. There are $4 \cdot 3^{n_m - 1}$ elements of F with distance n_m from e. Let c denote such an element. There are 3 branches at c in the Cayley graph of F that do not contain e. On each of these branches, choose a point a with $d(a, e) \geq \nu_m$. Repeat this process for each c with distance n_m from e. Let C_m be the collection of all the points a. Then $|C_m| = 4 \cdot 3^{n_m}$, and since $n_m \to \infty$ we know $|C_m| \to \infty$. For any distinct $a, b \in C_m$, the distance from e to $[a, b]$ is at most $n_m \leq \sqrt{\nu_m}$, because $[a, b]$ will always contain at least one of the points in F with distance n_m from e.

Since $|C_m| \to \infty$ and U is a non-principal ultrafilter on $\mathbb{N}$, we know that $\Pi_m C_m / U$ is a set of points in the asymptotic cone with cardinality 2^ω. Moreover, in the asymptotic cone we have

$$d([a_m], [e]) = \lim_U \frac{d_C(a_m, e)}{\nu_m} \geq 1.$$

Since $(a_m \cdot b_m)_e \leq \sqrt{\nu_m}$, we know the distance from $[e]$ to the geodesic segment connecting $[a_m]$ and $[b_m]$ is

$$([a_m] \cdot [b_m])_{[e]} = \lim_U \frac{(a_m \cdot b_m)_e}{\nu_m} = 0.$$

Thus, $[e] \in [[a_m], [b_m]]$, putting $[a_m]$ and $[b_m]$ on separate branches at $[e]$. This gives us 2^ω-many distinct branches at $[e]$ in $\mathrm{Con}_{U,(\nu_m)}(G, d_C, e)$. Therefore, there must be 2^ω-many branches at every point in the cone. $\square$

10.8 Corollary *Let G be a non-elementary hyperbolic group. Let $\mathcal{M}$ be the model of* rb$\mathbb{R}$T$_r$ *with underlying space equal to the closed r-ball of* $\mathrm{Con}_{U,(\nu_m)}(G, d_C, e)$*. Then* $\mathcal{M}$ *is the unique saturated model of* rb$\mathbb{R}$T$_r$ *of density* 2^ω*.*

Proof By the preceding lemma and Lemma 10.1, we know $\mathcal{M}$ has density 2^ω and $\mathcal{M}$ is 2^ω-saturated by Lemma 8.5. □

11 Models of rb$\mathbb{R}$T$_r$: Constructions and non-categoricity

In this section we show that rb$\mathbb{R}$T$_r$ has the maximum number of models of density character κ for every infinite cardinal κ. Indeed, for each κ we construct a family of 2^κ-many such models such that no two members of the family are homeomorphic. (Two models of rb$\mathbb{R}$T$_r$ are homeomorphic if their underlying $\mathbb{R}$-trees are homeomorphic by a map that takes base point to base point. Note that non-homeomorphic models of rb$\mathbb{R}$T$_r$ are necessarily non-isomorphic.) First we treat separable models, and the amalgamation techniques used in that case also allow us to characterize the principal types of rb$\mathbb{R}$T$_r$ and to show that this theory has no atomic model. Then we use simple amalgamation constructions to handle nonseparable models.

11.1 Lemma *Let S be a non-empty set of integers, each of which is ≥ 3. There exists a separable richly branching $\mathbb{R}$-tree M such that*

(1) *for each $k \in S$ the set $\{x \in M \mid x$ has degree $k\}$ is dense in M, and*

(2) *given a branch point $x \in M$ the degree of x is an element of S.*

Proof Let $(k_j \mid j \in \mathbb{N})$ be a sequence such that every element of S appears infinitely many times in the sequence, and every term of the sequence is an element of S. We construct an increasing sequence $N_0 \subseteq N_1 \subseteq \cdots \subseteq N_j \ldots$ of separable $\mathbb{R}$-trees as follows.

Let N_0 be the $\mathbb{R}$-tree $\mathbb{R}$ with base point 0. Let A_0 be a countable, dense subset of N_0. Use Lemma 6.1 to add $k_0 - 2$ distinct rays (copies of $\mathbb{R}^{\geq 0}$) at each point in A_0, bringing the number of branches of infinite length at each point in A_0 up to k_0. Call the resulting $\mathbb{R}$-tree N_1. Note that $N_0 \subseteq N_1$. The $\mathbb{R}$-tree N_1 is separable, since it is a countable union of separable spaces. Note also that all the points in $N_1 \setminus A_0$ only have 2 branches, and it is straightforward to show $N_1 \setminus A_0$ is uncountable and dense in N_1.

Once N_j has been constructed, to construct N_{j+1} let $A_j \subset N_j \setminus (\cup_{i=0}^{j-1} A_j)$ be a countable, dense subset of N_j. This is possible since $N_j \setminus (\cup_{i=0}^{j-1} A_j)$ is dense in N_j. Use Lemma

6.1 to add $k_j - 2$ rays at each point in A_j, bringing the number of branches of infinite length at each point in A_j up to k_j. The resulting $\mathbb{R}$-tree is N_{j+1}. Note that N_{j+1} is separable, since it is a countable union of separable spaces. Note also that all the points in $N_{j+1} \setminus \cup_{j=0}^{j} A_j$ still only have 2 branches, and that this set is uncountable and dense in N_{j+1}. Lastly, it is clear that given $x \in N_{j+1}$ the number of branches at x must be either 2 (in which case x is not a "branch point") or one of $\{k_0, \ldots, k_j\}$. This is because at the jth step, the only points at which we add rays are those in A_j, and then in subsequent steps we do not add rays at any of those points.

Let $M = \cup_{j \in \mathbb{N}} N_j$ be the union of this countable chain of separable $\mathbb{R}$-trees. Then M is a separable $\mathbb{R}$-tree (see Chiswell [8, Lemma 2.1.14].) Since for each N_j the number of branches at each branch point is an element of S, this will also be true in M. Let $k \in S$. Let $J(k) = \{j \in \mathbb{N} \mid k_j = k\}$. By how we chose the sequence (k_j) the set $J(k)$ is infinite. The set of points in M which have exactly k branches is $\cup_{j \in J(k)} A_j$. We will show this set is dense in M. Let $x \in M$. Let $j_x \in \mathbb{N}$ be the smallest integer such that $x \in N_{j_x}$. Let $j^* \in J(k)$ be such that $j^* > j_x$. We know that A_{j^*} is dense in N_{j^*}, and that $x \in N_{j_x} \subseteq N_{j^*}$. Therefore, there are points in $A_{j^*} \subseteq \cup_{j \in J(k)} A_j$ arbitrarily close to x. Our choice of $x \in M$ was arbitrary, therefore $\cup_{j \in J(k)} A_j$ is dense in M. Because we chose a non-empty S with members all ≥ 3 the set of branch points with at least 3 branches of infinite length is dense in M. Therefore, M is a richly branching $\mathbb{R}$-tree. $\square$

11.2 Remark   In the preceding proof, we did not use the fact that we are considering homeomorphisms of *pointed* topological spaces. The $\mathbb{R}$-trees constructed in Lemma 11.1 are in fact non-homeomorphic, even when we are not required to preserve the base point.

11.3 Theorem   *There exist 2^ω-many pairwise non-homeomorphic (hence non-isomorphic) separable models of* $rb\mathbb{R}T_r$.

Proof   Any homeomorphism g between models $\mathcal{M}$ and $\mathcal{N}$ of $rb\mathbb{R}T_r$ is a homeomorphism on the underlying $\mathbb{R}$-trees which must preserve branching. In particular, given $n \in \mathbb{N}^{\geq 3}$, if there is a point with degree n in M, then there must be a point with degree n in N. Choose two different subsets S and S' of $\mathbb{N}^{\geq 3}$, and construct a richly branching tree for each as in Lemma 11.1. Let $\mathcal{M}$ and $\mathcal{M}'$ be the models of $rb\mathbb{R}T_r$ based on the completions of their closed r-balls, respectively. By Lemma 2.18, taking the completion only adds points of degree 1. It follows that $\mathcal{M}$ and $\mathcal{M}'$ cannot be homeomorphic. Since there are 2^ω-many different such sets S, there are 2^ω-many different non-homeomorphic, separable models of $rb\mathbb{R}T_r$. $\square$

Recall that given a continuous theory T, a type $q \in S_n(T)$ is *principal* if for every model $\mathcal{M}$ of T, the set $q(\mathcal{M})$ of realizations of q in $\mathcal{M}$ is definable over the empty set. As in classical first order logic, given a complete theory in a countable signature, there is a Engeler-Ryll-Nardzewski-Svenonius theorem stating the equivalence between ω-categoricity and the fact that every type is principal. (See Ben Yaacov, Berenstein, Henson and Usvyatsov [4, Theorem 12.2].) Furthermore, a type q is principal if and only if q is realized in every model of T. (See [4, Theorem 12.6].)

Theorem 11.3 obviously implies that rb$\mathbb{R}$T$_r$ is not ω-categorical, and thus not every type of rb$\mathbb{R}$T$_r$ is principal. Our next result gives a characterization of the principal types in $S_n(\text{rb}\mathbb{R}\text{T}_r)$. In particular, a principal type is the type of a tuple of points that all lie along a single piecewise segment with p as an endpoint. Thus, there are very few of them. As a consequence, we conclude that rb$\mathbb{R}$T$_r$ does not have a prime model (equivalently, does not have an atomic model, one in which only principal types are realized).

For a clear and comprehensive treatment of separable models in continuous model theory, we refer the reader to Section 1 in Ben Yaacov and Usvyatsov [5]. Note that where we and Ben Yaacov, Berenstein, Henson and Usvyatsov [4] use the word *principal*, the authors of [5] use *isolated,* which is now the standard terminology. In ([5, Theorem 1.11]) they prove an omitting types theorem, and as a corollary ([5, Corollary 1.13]) show that nonprincipal types can be omitted. Further, it follows from [5, Definition 1.7] and properties of definable sets in continuous model theory, that every principal type is realized in every model, and this is implicit in the discussion following that definition.

11.4 Theorem *Let $q \in S_n(\text{rb}\mathbb{R}\text{T}_r)$. The following are equivalent:*

(1) *The type q is principal.*
(2) *There exists a permutation σ of $\{1, \ldots, n\}$ such that q contains the condition*

$$d(p, x_{\sigma(n)}) = d(p, x_{\sigma(1)}) + \sum_{i=1}^{n-1} d(x_{\sigma(i)}, x_{\sigma(i+1)}).$$

(3) *There exists a permutation σ of $\{1, \ldots, n\}$ such that for any model $\mathcal{M}$ of rb$\mathbb{R}$T$_r$ and $(b_1, \ldots, b_n) \in M^n$ that realizes q in $\mathcal{M}$, we have that $[p, b_{\sigma(1)}, \ldots, b_{\sigma(n)}]$ is a piecewise segment in $\mathcal{M}$.*

Proof That (2) and (3) are equivalent (for the same permutation σ) follows from Chiswell [8, Lemma 2.1.4]. See the discussion after Definition 2.6.

We show (1) implies (3) by proving the contrapositive. Let $\mathcal{M} \models \mathrm{rb}\mathbb{R}\mathrm{T}_r$ and suppose $(b_1, \ldots, b_n) \in M^n$ realizes q in $\mathcal{M}$. For each i let t_i be the number $d(p, b_i)$. Choose $j \in \{1, \ldots, n\}$ such that $d(p, b_j) \geq d(p, b_i)$ holds for every $i \in \{1, \ldots, n\}$. If (3) fails, there must be i such that b_i is not on the segment $[p, b_j]$. Note that this implies that p, b_i, b_j are distinct, so t_i and t_j are > 0. (If $d(p, p_j) = 0$ then $d(p, b_i) = 0$ for all i, meaning that $b_1 = \ldots b_n = p$, contradicting the assumption that (3) fails.)

Let $e = Y(p, b_i, b_j)$ be the closest point to b_i on the segment $[p, b_j]$. If $e = b_j$, then

$$d(p, b_i) = d(p, e) + d(e, b_i) = d(p, b_j) + d(b_j, b_i) > d(p, b_j)$$

which is a contradiction. This leaves two possibilities.

Case 1: Assume $e = p$. Then b_i and b_j are on different branches at p, so q contains the conditions $d(x_i, x_j) = d(p, x_i) + d(p, x_j)$, $d(p, x_i) = t_i$ and $d(p, x_j) = t_j$. It follows that whenever $(c_1, \ldots, c_n)$ realizes q in any model $\mathcal{N}$ of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$, the points c_i and c_j must be on different branches at p. However, using techniques as in 11.1, we can construct $\mathcal{N} \models \mathrm{rb}\mathbb{R}\mathrm{T}_r$ in which there is only a single branch at p. Then q is not realized in $\mathcal{N}$ and thus q is not principal.

Case 2: Assume $e \neq p$. Then p and b_j are on different branches at e, and Lemma 2.5 implies that $e \in [p, b_i]$ and $e \in [b_i, b_j]$. So p, b_i, b_j, being 3 distinct points, are on 3 distinct branches at $e = Y(p, b_i, b_j)$. As discussed in the preceding case, these properties of p, b_i, b_j are witnessed by conditions in q; also, q contains the condition $d(p, Y(p, x_i, x_j)) = t$ for some $t \in (0, r]$. It follows that in any model of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ in which q is realized, there must be a point c whose distance from p is t and at which the model has 3 branches. Using techniques discussed earlier in this section, we may build a model $\mathcal{N}$ in which there are no branch points at distance t from p. Then q is not realized in $\mathcal{N}$, implying that q is not a principal type.

Last, we show (3) implies (1). Let $q \in S_n(\mathrm{rb}\mathbb{R}\mathrm{T}_r)$ be a type satisfying the conditions in (3) (and hence in (2)) with permutation σ. Then $d(p, x_{\sigma(n)})^q \geq d(p, x_{\sigma(i)})^q$ for all $i \in \{1, .., n\}$. Take any model $\mathcal{M} \models \mathrm{rb}\mathbb{R}\mathrm{T}_r$. Then there is at least one branch at the base point p, and by Lemma 7.7 along this branch we can find a point b so that $d(p, b) = d(p, x_{\sigma(n)})^q$. On the segment $[p, b]$ we can find points $b_1, \ldots, b_n$ such that $d(p, b_{\sigma(i)}) = d(p, x_{\sigma(i)})^q$ for all i. From this it follows that $(b_1, \ldots, b_n)$ realizes q in $\mathcal{M}$. Since q is realized in every model of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$, it is a principal type. $\qquad \square$

11.5 Corollary   *The L-theory* $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ *has no prime model.*

Proof   Assume $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ has a prime model $\mathcal{M}$ with underlying $\mathbb{R}$-tree (M, d, p). Then $\mathcal{M}$ is atomic and the type of any tuple $b_1, \ldots, b_n$ must be principal. By the preceding

theorem, this means that for any pair b_1, b_2 in M, either $b_1 \in [p, b_2]$ or $b_2 \in [p, b_1]$. It follows that M consists of a piecewise segment with endpoint p. In particular, $\mathcal{M}$ is not richly branching, which is a contradiction. $\qquad\square$

We finish this section by showing that when κ is uncountable, then the number of different models of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ having density character equal to κ is also the maximum possible, namely 2^κ. As in the case $\kappa = \omega$, which was treated in the first part of this section, we produce large sets of models that are not only non-isomorphic, but in fact have underlying $\mathbb{R}$-trees which are non-homeomorphic (as pointed topological spaces).

We will carry out the construction by induction on κ, and we begin with a useful lemma.

11.6 Lemma *Let κ be an uncountable cardinal. The following conditions are equivalent:*

(1) *The number of non-homeomorphic models of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ of density character $\leq \kappa$ is at least κ.*

(2) *The number of non-homeomorphic models of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ of density character $\leq \kappa$ that have just one branch at the base point is at least κ.*

(3) *The number of non-homeomorphic models of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ of density character $= \kappa$ is 2^κ.*

Proof Let κ be an uncountable cardinal. Clearly, (3) implies (1). To show (2) implies (3), assume (2) and let $(\mathcal{B}_\alpha \mid \alpha < \kappa)$ be a list of the pairwise non-homeomorphic models of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$, each with density character $\leq \kappa$ and exactly one branch at the base point. Given a subset $S \subseteq \kappa$ of cardinality $= \kappa$, take the collection of B_α for $\alpha \in S$ and glue them all together at their base points using Theorem 6.4. Call this amalgam M_S, and let p be the point in M_S at which the B_α are all glued together. Make p the base point of the L_r-structure $\mathcal{M}_S$. Note that the density character of $\mathcal{M}_S$ is exactly κ, since each branch of its underlying $\mathbb{R}$-tree M_S at p has density $\leq \kappa$, and there are exactly κ-many branches at p. Moreover, it is easy to check that $\mathcal{M}_S$ is a model of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$.

By this construction, if B ranges over the branches of M_S at p, the homeomorphism type of $B \cup \{p\}$ (with p as distinguished element) ranges bijectively over the homeomorphism types of B_α (also with p as distinguished element) as α ranges over S. It follows that the homeomorphism type of $\mathcal{M}_S$ determines S. Therefore the family

$\{\mathcal{M}_S \mid S \subseteq \kappa$ and S has cardinality $= \kappa\}$ verifies condition (3), since κ has 2^κ-many subsets of cardinality $= \kappa$.

Finally, we prove that (1) implies (2). For each model $\mathcal{M}$ of rb$\mathbb{R}$T$_r$, let $b(M)$ denote the number of branches of M at its base point; we take this to be a positive integer or ∞, where $b(M) = \infty$ means that there are infinitely many branches. Since κ is uncountable, condition (1) yields a class $\mathcal{K}$ of at least κ-many non-homeomorphic models of rb$\mathbb{R}$T$_r$, each of density character $\leq \kappa$, such that $b(M)$ has a constant value b as M ranges over $\mathcal{K}$. We may assume $b \neq 1$, since otherwise condition (2) is satisfied by the models in $\mathcal{K}$.

Let a be an integer ≥ 3 that is different from $b + 1$. (Note that $a \neq b + 1$ is automatically true if b is ∞.) Using a method similar to that in the proof of Lemma 11.1, we may take $\mathcal{N} = (N, d, p)$ to be a separable model of rb$\mathbb{R}$T$_r$ with the following properties: (1) for all x in N, the number of branches in N at x is 1 or 2 or a; (2) N has a single branch at its base point p; and (3) N has a single branch at some point y, where $d^{\mathcal{N}}(p, y) = r/2$. To get single branches at 2 points with a given distance as required here, proceed as in the proof of Lemma 11.1, starting with the interval $[0, r/2] \subseteq \mathbb{R}$ instead of $\mathbb{R}$, and in subsequent steps always exclude 0 and $r/2$ from the sets of points where rays are added.

Now consider an arbitrary $\mathcal{M} \in \mathcal{K}$, and denote the base point of $\mathcal{M}$ by q. Scale the metric on M down by a factor of 2, resulting in an L_r-structure with a radius of $r/2$. We construct a larger $\mathbb{R}$-tree M^* by amalgamating the scaled-down M and $\mathcal{N}$ in the way that identifies q and y; we will denote this point of M^* by qy. We take the base point of $\mathcal{M}^*$ to be the base point p of $\mathcal{N}$. The radius of $\mathcal{M}^*$ is r, because the radius of $\mathcal{N}$ was r, and by our amalgamation construction, for every point $x \in M^*$ residing in the scaled down copy of M,

$$d^{\mathcal{M}^*}(p, x) = d^{\mathcal{M}^*}(p, qy) + d^{\mathcal{M}^*}(qy, x) = \frac{r}{2} + \frac{d^{\mathcal{M}}(y, x)}{2}$$

which attains its maximum value r as x ranges over the scaled down copy of M.

It is straightforward to check that $\mathcal{M}^*$ is a model of rb$\mathbb{R}$T$_r$, has density $\leq \kappa$, and has a single branch at its base point. Note that the branches of M^* at the amalgamated point qy consist of the branches of q in M together with the tree that results from N by removing y. In particular, this means that M^* has $b + 1$-many branches at qy.

We claim that the class $\mathcal{K}^* = \{\mathcal{M}^* \mid \mathcal{M} \in \mathcal{K}\}$ verifies condition (3); it remains only to show that no two members of this class are homeomorphic. (Recall that the homeomorphisms we consider must take base point to base point.) The key to this is

the fact that the point qy can be topologically identified in M^*, given that we know the base point p. To do this, note first that the segment $X = [p, qy)$ in M^* is identical to the segment $[p, y)$ in N, and every point in X has the same number of branches in M^* as in N. Therefore every point x of X has 1, 2, or a-many branches in M^*, and thus the number of branches at x is different from the number of branches at qy. From this we conclude that for any $\mathcal{M}_1, \mathcal{M}_2 \in \mathcal{K}$, any homeomorphism of M_1^* onto M_2^* that takes base point to base point must map the scaled version of M_1 onto the scaled version of M_2. Since this can only happen when $\mathcal{M}_1 = \mathcal{M}_2$, by assumption on $\mathcal{K}$, we conclude that $\mathcal{M}_1^* = \mathcal{M}_2^*$, as desired. $\qquad\square$

11.7 Theorem *Let κ be an uncountable cardinal. The number of non-homeomorphic models of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ of density character equal to κ is 2^κ.*

Proof We assume that σ is the least uncountable cardinal at which there are strictly fewer than 2^σ-many non-homeomorphic models of density character equal to σ, and derive a contradiction. Using Theorem 11.3, we see that condition (1) in Theorem 11.6 holds when $\kappa = \omega_1$; condition (3) in that result yields that $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ has 2^{ω_1}-many non-homeomorphic models of density character equal to ω_1. Thus $\sigma > \omega_1$. Now suppose σ is a successor cardinal; say it is the next cardinal bigger than λ, which must be uncountable. Our choice of σ ensures that there must be $2^\lambda \geq \sigma$-many non-homeomorphic models of density character λ. Applying Lemma 11.6 with $\kappa = \sigma$ gives a contradiction; indeed, we have verified condition (1), while condition (3) is false. So σ must be a limit cardinal. Let τ be the number of non-homeomorphic models of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ that have density character $\leq \sigma$; our treatment of ω_1 shows that τ is uncountable. Furthermore, Lemma 11.6 applied to $\kappa = \sigma$ yields $\tau < \sigma$. Our choice of σ ensures that there are $2^\tau > \tau$-many non-homeomorphic models of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ that have density character τ, contradicting the definition of τ. $\qquad\square$

12 Unbounded $\mathbb{R}$-trees

As noted in the Introduction, we have chosen to treat *bounded* pointed $\mathbb{R}$-trees in this paper, because many of the model-theoretic ideas and tools we need from continuous first order logic are only documented for bounded metric structures in the literature.

However, it would certainly be a natural research topic to study the model theory of unbounded (ie, not necessarily bounded) pointed $\mathbb{R}$-trees. The most immediately available setting for doing this would be to consider a pointed $\mathbb{R}$-tree (M, p) as a

many-sorted metric structure in which each sort is one of the (closed) bounded balls of (M, p) (centered at p), and the union of the family of distinguished balls is all of M. Everything done in this paper can easily be carried over to that setting. The disadvantages of doing so are the technical awkwardness of the many-sorted framework and the need for imposing an arbitrary family of radii for the bounded balls into which the full tree is stratified.

It is certainly more mathematically natural to consider pointed $\mathbb{R}$-trees on their own, without imposing a many-sorted stratification. There are suitable logics for doing model theory with such unbounded structures. For example, a version of continuous first order logic for unbounded metric structures is described in Ben Yaacov [1]. Also, a logic based on *positive bounded formulas* and an associated concept of *approximate satisfaction* is presented in Section 6 of Dueñez and Iovino [11]. However, for neither of these approaches are the ideas and tools of model theory developed as we need them in this paper.

In each of these three available settings for treating arbitrary pointed $\mathbb{R}$-trees, the arguments in this paper can be used easily to demonstrate: (1) the class of pointed $\mathbb{R}$-trees is axiomatizable and (2) for each $r > 0$, the ball $\{x \mid d(x, p) \leq r\}$ is a definable set (over $\emptyset$, uniformly in all pointed $\mathbb{R}$-trees). Together with what is developed in Ben Yaacov, Berenstein, Henson and Usvyatsov [4] as well as in Ben Yaacov [1] and in Dueñez and Iovino [11], this quickly yields that the model theoretic frameworks for pointed $\mathbb{R}$-trees provided by these three settings are completely equivalent. In particular, this approach yields a model companion for the theory of pointed $\mathbb{R}$-trees whose models are exactly the richly branching $\mathbb{R}$-trees (ie, the complete pointed $\mathbb{R}$-trees described in Remark 7.2). Furthermore, this model companion has suitably stated versions of all the properties of $\mathrm{rb}\mathbb{R}\mathrm{T}_r$ that are proved in this paper.

References

[1] **I Ben Yaacov**, *Continuous first order logic for unbounded metric structures*, J. Math. Logic 8 (2008) 197–223; https://dx.doi.org/10.1142/S0219061308000737

[2] **I Ben Yaacov** *Modular functionals and perturbations of Nakano spaces*, J. Logic and Analysis 1 (2009) paper 1, 42 pages; https://dx.doi.org/10.4115/jla.2009.1.1

[3] **I Ben Yaacov**, **A Berenstein**, **C W Henson**, *Model-theoretic independence in the Banach lattices $L_p(\mu)$*, Isr. J. Math. 183 (2011) 285–320; https://dx.doi.org/10.1007/s11856-011-0050-4

[4] **I Ben Yaacov**, **A Berenstein**, **C W Henson**, **A Usvyatsov**, *Model theory for metric structures*, in "Model Theory with Applications to Algebra and Analysis, Vol II", (eds Z Chatzidakis, D Macpherson, A Pillay, and A Wilkie), Lecture Notes series of the London Mathematical Society No 350, Cambridge University Press (2008) 315–427; https://dx.doi.org/10.1017/CBO9780511735219.011

[5] **I Ben Yaacov**, **A Usvyatsov**, *On d-finiteness in continuous structures*, Fund. Math. 194 (2007) 67–88; https://dx.doi.org/10.4064/fm194-1-4

[6] **I Ben Yaacov**, **A Usvyatsov**, *Continuous first order logic and local stability*, Trans. Amer. Math. Soc. 362 (2010) 5213–5259; https://dx.doi.org/10.4115/10.1090/S0002-9947-10-04837-3

[7] **M R Bridson**, **A Haefliger**, *Metric Spaces of Non-Positive Curvature*, Grundlehren der mathematischen Wissenschaften 319 Springer-Verlag (1999); https://dx.doi.org/10.1007/978-3-662-12494-9

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

[9] **L van den Dries**, **A Wilkie**, *Gromov's theorem on groups of polynomial growth and elementary logic*, J. Algebra 89 (1984) 349–374; https://dx.doi.org/10.1016/0021-8693(84)90223-0

[10] **C Drutu**, *Quasi-isometry invariants and asymptotic cones*, Int J Algebra and Computation 12 (2002) 99–135; https://dx.doi.org/10.1142/S0218196702000948

[11] **E Dueñez**, **J Iovino**, *Model theory and metric convergence I: Metastability and dominated convergence*, in "Beyond First Order Model Theory", (ed J Iovino), CRC Press (2017) 131–187; https://dx.doi.org/10.1201/9781315368078

[12] **A G Dyubina**, **I Polterovich**, *The structure of hyperbolic spaces at infinity* (Russian), Uspekhi Mat. Nauk 53 (1998) 239–240; translation in Russian Math. Surveys 53 (1998) 1093–1094; https://dx.doi.org/10.1070/rm1998v053n05ABEH000080

[13] **A G Dyubina**, **I Polterovich**, *Explicit constructions of universal $\mathbb{R}$-trees and asymptotic geometry of hyperbolic spaces*, Bull. London Math. Soc. 33 (2001) 727–734; https://dx.doi.org/10.1112/S002460930100844X

[14] **W Hodges**, *Model Theory*, Cambridge University Press (1993); https://dx.doi.org/10.1017/CBO9780511551574

[15] **J Roe**, *Lectures on Coarse Geometry*, University Lecture Series 31, Amer. Math. Soc. (2003); https://dx.doi.org/10.1090/ulect/031

Sylvia Carlisle, Rose-Hulman Institute of Technology, Terre Haute, Indiana 47803, USA

C. Ward Henson, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA

carlisle@rose-hulman.edu

https://faculty.math.illinois.edu/~henson