UNBOUNDED CONTINUOUS LOGIC AND METRIC GEOMETRY

UNBOUNDED CONTINUOUS LOGIC AND METRIC GEOMETRY

By

MATTHEW LUTHER, B.A., M.SC.

A Thesis

Submitted to the School of Graduate Studies

in Partial Fulfillment of the Requirements

for the Degree

Doctor of Philosophy

McMaster University

DOCTOR OF PHILOSOPHY (2015)	McMaster University
(Mathematics)	Hamilton, Ontario

TITLE:	Unbounded Continuous Logic and Metric Geometry
AUTHOR:	Matthew Luther, B.A. (The College of New Jersey), M.Sc. (McMaster University)
SUPERVISOR:	Professor Bradd Hart
NUMBER OF PAGES:	vi, 154

Abstract

This thesis studies continuous logic and its application to metric geometry. An adaptation of continuous logic for unbounded pointed metric spaces is introduced and developed. Background on CAT(k) spaces, asymptotic cones, symmetric spaces, and buildings is provided. Various definability results are proved regarding geodesic rays and the building structure on them. We conclude with a proof of the instability of asymptotic cones of a certain class of symmetric spaces.

Acknowledgements

I'd like to thank my supervisor Dr. Bradd Hart. Over the past few years he has been generous with both his support and his espresso. His guidance and willingness to endure long and frequent conversations were invaluable throughout this work.

I'd also like to thank my committee members Dr. Deirdre Haskell and Dr. Patrick Speissegger, along with Dr. Matt Valeriote. I've appreciated the activeness of the logic group at McMaster.

In the last summer of this thesis, I ran a seminar series on the model theoretic contents. I thank the attendees of this seminar.

Finally, I thank the friends I've made in the math department, my former housemates, my family, and my girlfriend Celina. They have all been supportive over the years.

Contents

1. Introduction

This thesis is divided into two parts. The first part introduces and develops a version of continuous logic and model theory for unbounded pointed metric structures. The second part uses this logic to study important spaces and constructions in metric geometry.

A version of continuous logic for bounded metric structures was introduced in [3] and has earlier roots. It is a $[0,1]$-valued logic which uses bounded metric spaces as sorts, and lets the structure carry bounded, uniformly continuous functions and relations. The quantifiers in this setting are given by sup and inf, rather than the $\forall$ and $\exists$ quantifiers of discrete first-order logic. Many basic results from discrete model theory have analogs in this setting. So, it has a strong claim to the title of first order logic for metric structures.

However, many important metric spaces are unbounded, and dealing with unbounded spaces in this logic is handled in ad hoc or undesirable ways. Usually, one has to think of the space as an infinite family of nested, bounded sorts, or else have some reason to know that a single bounded subset is sufficient to understand the space. Including the balls as sorts sometimes imposes unwelcome structure. The restriction on functions and relations can also be awkward to work around. For example, addition on the real line would have to be viewed as an infinite family of addition functions from balls to larger balls.

In [4], Ben Yaacov addresses these shortcomings. There, he presents a continuous logic for unbounded metric structures and applies it to study perturbations of norms. Elsewhere, he uses it to study metric valued fields. Both areas of study are awkward to work with in the bounded logic. A fundamental issue in the unbounded setting is the role played by quantifiers in the logic. Ben Yaacov handles this by restricting the application of sup and inf quantifiers to a subclass of formulas which behaves well.

The version of continuous logic developed in this thesis takes a different approach. We will focus on unbounded pointed metric spaces. We introduce new quantifiers to account for the problems with sup and inf in the unbounded setting. We also allow a wider class of functions and formulas in our structures and logic. Ultimately, these functions and formulas are implicitly in the other versions of continuous logic, but we emphasize their status from the beginning. This is motivated by the application in the second part of the thesis. There, we are almost entirely concerned with quantifation over sequences in an unbounded space. Our development fundamentally acknowledges this possibility. We feel this approach provides more agreeable semantics.

The novel content in this development is the definition of this wide class of functions and formulas, the new quantifiers, working with sequences, and demonstrating the analogs of basic model theoretic results in this setting. Many of the basic results can be carried over from the bounded case, or adapted with little trouble. But there are also many subtleties introduced in some fundamental notions, most notably in the sections on ultraproducts and Łoś's Theorem and on definable sets.

The second part of this thesis studies CAT(κ) spaces and asymptotic cones of symmetric spaces of noncompact type. These are fundamental objects in metric geometry. We review some background material, and then propose a model theoretic approach to the area. Various definability results are proved, and instability of these asymptotic cones is demonstrated.

CAT(κ) spaces M are spaces with a notion of bounded curvature described via the comparison of triangles in M with triangles in euclidean spaces, hyperbolic spaces, and spheres. The curvature bounds in these spaces makes them very amenable to geometric arguments that work in the more familiar spaces. We provide axiomatizations of CAT(κ) spaces in our logic, and prove several definability results for geodesics in CAT(κ) spaces. For example, we show how in a language with just the distance predicate, we can obtain the set of geodesic rays in a CAT(0) space as a definable set of sequences.

Asymptotic cones were introduced by Gromov in his work [13] characterizing groups of polynomial growth. A definition later in terms of ultrafilters was given by van den Dries and Wilkie in [19]. The concept has seen ongoing interest in the decades since. A main reference for us throughout the second half is the work of Kleiner and Leeb [14], where asymptotic cones of symmetric spaces are studied as a way to understand the boundary of the space and prove an important rigidity result. Asymptotic cones are a construction where one takes a metric space (X, d) and produces a limit space out of the rescaled metrics $(X, \frac{d}{n})$. This is sometimes described intuitively as taking the tangent space at infinity, as opposed to the tangent space at a point, which would be by seen by something like $(X, n \cdot d)$. We check that the asymptotic cone construction is an ultraproduct in our setting, and prove a few results about the general construction.

Our interest and work on the cones of symmetric spaces was inspired by the paper [15] and its main result, which we now state.

Theorem 1.0.1. *Suppose G is a connected semisimple Lie group with at least one absolutely simple factor S with $\mathbb{R}$-rank$(S) \geq 2$ and let Γ be a uniform lattice in G.*

If the continuum hypothesis holds, then Γ has a unique asymptotic cone up to homeomorphism.

If the continuum hypothesis fails, then Γ has $2^{2^{\aleph_0}}$ many asymptotic cones up to homeomorphism.

This result addresses a question posed by Gromov about when finitely generated or finitely presented groups have unique asymptotic cones. To a model theorist having familiarity with continuous logic, the result and its proof suggest an instability result. Namely, instability is suggested by the combination of Theorem 5.6 of [12], a result from [11] which we include as Proposition 4.0.9, and the argument in [15] constructing different ultrafilters to distinguish order types of a field associated to the asymptotic cones. The first of these in particular proves the following.

Theorem 1.0.2. *If A is a separable metric structure in a separable language and the theory of A is stable, then the ultrapowers $A^{\mathcal{U}}$ and $A^{\mathcal{V}}$ are isomorphic for any two non-principal ultrafilters $\mathcal{U}$, $\mathcal{V}$ on $\mathbb{N}$.*

The proof of this result involves showing that the ultrapowers are saturated (and elementary equivalent since they are both ultrapowers of A), and so must be isomorphic. In our setting, Bradd Hart suggested an analogous argument that if the asymptotic cones of Γ were stable for all ultrafilters $\mathcal{U}$, then even when the continuum hypothesis fails, one would expect elementary equivalence to imply isomorphism, and so there would be at most $2^{\aleph_0}$ many asymptotic cones. Since the result above finds $2^{2^{\aleph_0}}$ many asymptotic cones, this points to at least one asymptotic cone being unstable. The argument in [15] then suggests that the instability should come from an ordering in the field associated to the asymptotic cone.

We prove this instability result in the final section of the thesis, Section 11.2. Our general strategy is to interpret the associated field in the theory of each asymptotic cone, and to demonstrate the order property within the field.

A substantial amount of the second part of this thesis builds up the background to discuss the objects involved. We need to know enough about symmetric spaces, euclidean and spherical buildings, and projective geometry. We prove definability of geodesics, certain subsets of geodesic rays which capture the spherical building structure in asymptotic cones of a symmetric space, and finally obtain an associated projective plane and the desired field.

All of these definability results develop a foundation for studying these spaces via continuous logic. We suspect the meeting of continuous logic with these spaces has value because of the style of many arguments and importance of ultraproducts in metric geometry. For example, once familiar with our logic, one can see definability results and special cases of Łoś's Theorem in [14].

Part 1: Model Theory

The first part of this thesis develops a version of continuous logic. The approach is novel, but draws heavily from [5], [3], and [4]. This particular version is designed for dealing with unbounded, pointed spaces. The guiding motivation is to obtain a framework for working with first-order-definable sequences in ultraproducts of unbounded metric spaces.

2. Unbounded Continuous Logic

Continuous logic is a real-valued logic that fundamentally depends on the notion of distance rather than equality. The objects it applies to are complete, pointed metric spaces endowed with functions. But, in order to develop a useful theory, we have to restrict our attention somewhat. The reason is that we want to be able to carry out certain limit constructions, and not all functions behave well under limits.

2.1. **A collection of guiding examples.** We begin with a few simple examples to guide the construction of our logic. In each example, we take a sequence of metric spaces M_n and functions and consider candidates for the limit spaces M and limit functions. The discussion here is informal and is meant to develop awareness of key features.

Example 2.1.1. Let $M_n = \{0, n\}$ with the usual metric $d(0, n) = n$. There are two natural candidates for the limit space M:

- One has a point at infinity, $M = \{0, \infty\}$ where $d(0, \infty) = \infty$, but this is not a metric space.
- The other is $M = \{0\}$, which discards the point at infinity and remains a metric space.

Example 2.1.2. Let $M_n = \{0, \frac{1}{n}\}$ with the usual metric, and let f_n be defined by $f_n(0) = 0$ and $f_n(\frac{1}{n}) = 1$.

- We could let the limit consist of two distinct points, that is $M = \{0, 0'\}$ with $d(0, 0') = 0$, and let the limit function be defined by $f(0) = 0$ and $f(0') = 1$. This is not a metric space.
- If M should be a metric space, it should just be $M = \{0\}$. In other words, the sequences $(0)_{n\in\mathbb{N}}$ and $(\frac{1}{n})_{n\in\mathbb{N}}$ should be identified. But then $f = \lim f_n$ does not have a clear definition.

Example 2.1.3. Let $M_n = [0, 1]$ with the usual metric, and f_n given by $f_n(x) = x^n$. The limit should be $M = [0, 1]$ with the same metric. In this case, there is also a seemingly clear candidate for the limit function, namely the pointwise limit of the f_n, which is f with $f(x) = 0$ for $x < 1$ and $f(1) = 1$.

However, f not being continuous leads to some ambiguity. Consider the sequence $(1 - \frac{1}{n})_{n\in\mathbb{N}}$ viewed as an element of $\prod_{n\in\mathbb{N}} M_n$. It is reasonable to say that this sequence tends to $1 \in M$. But $f(1) = 1$, while $\lim_{n\to\infty} f_n(1 - \frac{1}{n}) = e^{-1} \neq 1$.

Example 2.1.4. Let $M_n = \{0\}$ and define functions $f_n : M_n \to \mathbb{R}$ by setting $f_n(0) = n$. The expected limit space $M = \lim_n M_n$ is $\{0\}$, but $\lim f_n(0)$ tends to infinity, a value outside of $\mathbb{R}$.

Example 2.1.5. Let M_n be the set $\{0, 1 - \frac{1}{n}\}$ with the usual distance. We would expect $M = \{0, 1\}$ as the limit space. Notice that $\sup_{x,y \in M_n} d(x, y) = 1 - \frac{1}{n}$, and that this sequence has limit 1, which agrees with $\sup_{x,y \in M} d(x, y) = 1$.

However, if we instead consider supremums $\sup_{x,y \in B_1(M_n)} d(x, y)$ over the open ball of radius 1 rather than the entire space (in this case, the closed ball of radius 1 is the whole space), the sequence is still $1 - \frac{1}{n}$ and has limit 1, but the open ball of radius 1 in M only has a single point and so $\sup_{x,y \in B_1(M)} = 0$.

Example 2.1.6. Let M_n be the set $\{0, 1 + \frac{1}{n}\}$ with the usual distance. Again, the limit space should be $M = \{0, 1\}$. Similarly to the last example, there is a supremum which is poorly behaved with respect to the limit, but this time it is the supremum over the closed ball of radius 1 which is problematic.

Example 2.1.7. For each $n \in \mathbb{N}$, let M_n be $\{0, 1\}^{\mathbb{N}}$, and define f_n by

$$f_n(x) = \max_{k \leq n}(x_k)$$

for each $x \in M_n$. For each n, define two elements $x, y \in M_n$ which are themselves sequences $(x_{n,k} : k \in \mathbb{N})$ and $(y_{n,k} : k \in \mathbb{N})$ as follows.

$$x_{n,k} = \begin{cases} 0 & \text{when } k < n \\ 0 & \text{when } k = n \\ 1 & \text{when } k > n \end{cases} \qquad y_{n,k} = \begin{cases} 0 & \text{when } k < n \\ 1 & \text{when } k = n \\ 1 & \text{when } k > n \end{cases}$$

Then it would be reasonable to say that x and y both converge (as $n \to \infty$) to the 0-sequence. But $f_n(x_n) = 0$ for all n, while $f_n(y_n) = 1$ for all n. So these sequences always differ by 1 in some coordinate.

2.2. **Basic definitions and notation.** This section collects a few general definitions and notation we will use throughout our discussions.

Notation 2.2.1. We will indicate when a tuple or sequence has finite length by referring to it as a finite tuple or finite sequence. Otherwise, it is not assumed to be constrained and could have any index set. That is, a tuple or sequence might be finite or infinite, and not necessarily indexed by $\mathbb{N}$ when infinite. Though, we will generally prefer tuple for finite lengths, and prefer sequence for arbitrary lengths.

Notation 2.2.2. We will generally write tuples or sequences with notation like $(x_n : n \in \alpha)$ or $(x_n)_{n \in \alpha}$. Once a sequence is introduced, we might drop the index set from the notation but emphasize that it is still a sequence rather than a coordinate by maintaining the parentheses, as in (x_n). Often we will use the unsubscripted letter as a name for the sequence, as in $x = (x_n : n \in \alpha)$.

For convenience, we might introduce the sequence just as a single letter, such as writing $x \in \mathbb{R}^\alpha$, but refer to coordinates of x by using subscripts, essentially assuming that we have written $x = (x_n : n \in \alpha)$.

Notation 2.2.3. When we want to call attention to the index set, say the index set I for the sequence $x = (x_i : i \in I)$, we will refer to this as an I-indexed sequence. If we are discussing a projection of this sequence onto coordinates $J \subseteq I$, we might refer to this as the J-projection of $(x_i : i \in I)$, or denote it by $\pi_J(x)$.

This thesis is concerned almost entirely with the following spaces.

Definition 2.2.4. A pointed metric space $(M, d, \star)$ is a metric space (M, d) with a distinguished basepoint $\star \in M$.

Notation 2.2.5. We will write $\mathbb{R}_+$ for the positive reals, and similarly $\mathbb{N}_+$ for the positive subset of the naturals $\mathbb{N} = \{0, 1, 2, \dots\}$. We write $\mathbb{R}_{\geq 0}$ for nonnegative reals. The default metrics assumed for $\mathbb{R}$ and $\mathbb{N}$ are the usual $d(x, y) = |x - y|$. The default basepoint for $\mathbb{R}$ is taken to be 0.

Geometers have identified the following subclass of metric spaces, and unsurprisingly it plays an interesting role in continuous logic.

Definition 2.2.6. A metric space is **proper** if every closed ball is compact.

Proper spaces are to continuous logic what finite sets are to discrete logic. The meaning of this last statement becomes clear once the fundamental results in model theory are established in this setting.

We will often be interested in products of pointed metric spaces. To simplify the discussion, we will assume the following default basepoint and metric for finite products.

Definition 2.2.7. Let α be a finite set and let $(M_n, d_n, \star_n)$ for $n \in \alpha$ be pointed metric spaces. Our default basepoint for $\prod M_n$ will be taken to be $(\star_n : n \in \alpha)$, and our default metric for $\prod M_n$ will be given by defining the distance between $(x_n : n \in \alpha)$ and $(y_n : n \in \alpha)$ to be the maximum over the distances taken coordinate-wise, that is,

$$\max_{n \in \alpha} \{d_n(x_n, y_n) : n \in \alpha\}.$$

In any pointed metric space, there is a natural notion of magnitude for points given by $d(x, \star)$. We will often need to refer to this in our discussions, both for single and product spaces, so we use the following notation.

Notation 2.2.8. We will write $||x||$ for the distance $d(x, \star)$ and call this the magnitude of x. This notation extends to the case where x is a point in a product of pointed metrics and we use the default basepoint and metric above. Also, instead of writing $||x|| \leq r$ with $r \in \mathbb{R}$, we might just say that x is bounded by r.

So, for example, saying $||x|| \leq r$ in a product space with $r \in \mathbb{R}$ means that every coordinate x_n of x satisfies $d_n(x_n, \star_n) \leq r$.

We will sometimes need to generalize this last situation. For example, if the first coordinate needs to be bounded by r_1 but the second coordinate bounded by r_2. We will handle this with a generalized notation for balls around the basepoint.

Notation 2.2.9. Let $(M_n, d_n, \star_n)$ be pointed metric spaces for $n \in \alpha$, and let $M = \prod M_n$ with the default metric and basepoint. Let $r \in \mathbb{R}^\alpha_{\geq 0}$ be a sequence of nonnegative reals. We will write $B_r(M)$ to mean the set of $x \in M$ such that $||x_n|| < r_n$ for all $n \in \alpha$. This means that

$$B_r(M) = \prod_{n \in \alpha} B_{r_n}(M_n)$$

where for each $n \in \alpha$, the set $B_{r_n}(M_n)$ is the open d_n-ball of radius r_n centered at $\star_n$ in M_n. We will similarly handle closed balls, using $\bar{B}$ rather than B to indicate that it is closed.

If we need to emphasize the metric or basepoint, we will add them inside the parentheses, as in $B_r(M, d)$, $B_r(M, \star)$, or $B_r(M, d, \star)$.

2.3. **Controlled functions.** In our logic, we will allow our structures to carry certain associated functions. This section defines the class of functions we allow. These functions generalize the behavior of continuous functions on proper spaces. For example, all continuous functions on $\mathbb{R}$ will be in this class.

The most notable properties of these functions are that, when restricted to any ball around the basepoint, they are uniformly continuous and bounded. This makes the functions work well with the ultraproduct construction that we discuss in 3.7.

It will turn out that to obtain expected model theory results, we need to discuss certain limits of functions with those properties. This should be familiar to readers acquainted with bounded continuous logic, where there is a distinction between formulas and definable predicates. To avoid this distinction, we will use a larger class of functions from the start.

The motivation for the definition below is most evident in the ultraproduct construction and the proof of the analog of Łoś's Theorem in this setting (Theorem 3.7.4). This class of functions seems to arise straightforwardly from attempting to assume as little as possible in the proof of Łoś's Theorem. Notably, we will allow functions on arbitrarily indexed products in our logic. The applications in the second part of this thesis make this seem like a natural admission.

Definition 2.3.1. Let α be any set, and let $(M_n, d_n, \star_n)$ for $n \in \alpha$ be pointed metric spaces. Let $(M, d, \star)$ be any pointed metric space, and let

$$f : \prod_{n \in \alpha} M_n \to M$$

be a function.

We say f is **controlled** if there are functions

- $\lambda : \mathbb{R}_+^\alpha \to \mathbb{R}_+$
- $N : \mathbb{R}_+^\alpha \times \mathbb{R}_+ \to \mathcal{P}_{\text{fin}}(\alpha)$
- $\delta : \mathbb{R}_+^\alpha \times \mathbb{R}_+ \to \mathbb{R}_+$

such that for all $r \in \mathbb{R}_+^\alpha$ and $\varepsilon > 0$, and for all $x \in \prod_{n\in\alpha} M_n$ and $y \in \prod_{n\in\alpha} M_n$, the following hold.

(1) If for all $n \in N(r,\varepsilon)$ we have $||x_n|| < r_n$, then $||f(x)|| \leq \lambda(r)$.
(2) If for all $n \in N(r,\varepsilon)$ we have $||x_n|| < r_n$, $||y_n|| < r_n$, and $d(x_n, y_n) < \delta(r,\varepsilon)$, then $d(f(x), f(y)) \leq \varepsilon$.

In this case, we say that f is **controlled by** (λ, N, δ), and that (λ, N, δ) are **controllers** or **controlling functions** for f. Note that controllers for f are not unique.

We will refer to any triple of functions (λ, N, δ) as **controllers** if they are controllers for some f.

Before we check some properties of controlled functions, we will make a comment on the roles of the functions λ, N, and δ. Together, these functions answer the question of what needs to be known about the inputs to f in order to get some property, either a bound or some amount of continuity. If we provide a sequence of bounds $r \in \mathbb{R}_+^\alpha$ we are willing to admit for the coordinates of the inputs and a maximum allowable $\varepsilon > 0$, then N tells us what finite subset of coordinates n we must check are bounded and close enough, δ tells us how close is close enough for these coordinates, and λ tells us how large the output can possibly be.

Proposition 2.3.2. *Compositions of controlled functions are controlled.*

Proof. Let t be a function of the form $f(t_k : k \in \alpha)$ where f is controlled by $(\lambda_f, N_f, \delta_f)$ and each $t_k^M : \prod_{n\in\alpha_k} M_{(k,n)} \to M_k$ is controlled by $(\lambda_k, N_k, \delta_k)$, respectively. The domain of t is $\beta = \bigcup_{k\in\alpha} (\{k\} \times \alpha_k)$. We make the following definitions of the controllers for t.

- For all $r = (r_{(k,n)}) \in \mathbb{R}_+^\beta$, define $r^* = (r_k^* : k \in \alpha)$ by
$$r_k^* = \lambda_k(r_{(k,n)} : n \in \alpha_k) + 1$$
for each $k \in \alpha$.
- Define λ by $r \mapsto \lambda_f(r^*)$.
- For each r as above and $\varepsilon > 0$, let $N_{r,\varepsilon} \in \beta$ denote the finite set of pairs (k,n) such that $k \in N_f(r^*,\varepsilon)$ and $n \in N_k((r_{(k,m)} : m \in \alpha_k), \varepsilon)$.
- Define N by $(r,\varepsilon) \mapsto N_{r,\varepsilon}$.
- Define δ by sending (r,ε) to the minimal value of
$$\delta_k((r_{(k,n)} : n \in \alpha_k), \frac{1}{2}\delta_f(r^*,\varepsilon))$$
with $k \in N_f(r^*,\varepsilon)$.

Let $r = (r_{(k,n)}) \in \mathbb{R}_+^\beta$ and $\varepsilon > 0$.

Suppose $||x_{(k,n)}|| < r_{(k,n)}$ for all $(k,n) \in N(r,\varepsilon)$. Then for all $k \in N_f(r^*,\varepsilon)$ and all n in $N_k((r_{(k,m)} : m \in \alpha_k), \varepsilon)$, we have $||x_{(k,n)}|| < r_{(k,n)}$. Hence, for all $k \in N_f(r^*,\varepsilon)$, we have

$$||t_k^M(x)|| \leq \lambda_k(r_{(k,n)} : n \in \alpha_k) < r_k^*$$

and thus

$$||f^M(t_k^M(x) : k \in \alpha)|| \leq \lambda_f(r^*) = \lambda(r).$$

Suppose for all $n \in N(r,\varepsilon)$ that we have

- $||x_{(k,n)}|| < r_{(k,n)}$,
- $||y_{(k,n)}|| < r_{(k,n)}$, and
- $d(x_{(k,n)}, y_{(k,n)}) < \delta(r,\varepsilon)$.

Then for all $k \in N_f(r^*,\varepsilon)$ and all n in $N_k((r_{(k,m)} : m \in \alpha_k), \varepsilon)$, we have $||x_{(k,n)}|| < r_{(k,n)}$, $||y_{(k,n)}|| < r_{(k,n)}$, and

$$d(x_{(k,n)}, y_{(k,n)}) < \delta_k((r_{(k,n)} : n \in \alpha_k), \frac{1}{2}\delta_f(r^*,\varepsilon))$$

Hence, for all $k \in N_f(r^*,\varepsilon)$, we have $||t_k^M(x)|| < r_k^*$, $||t_k^M(y)|| < r_k^*$, and

$$d(t_k^M(x), t_k^M(y)) < \delta_f(r^*,\varepsilon).$$

This implies

$$d(f^M(t_k^M(x) : k \in \alpha), f^M(t_k^M(y) : k \in \alpha)) \leq \varepsilon.$$

$\square$

The infinitary controlled functions always arise as certain kinds of limits of the finitary controlled functions.

Definition 2.3.3. Let $f : \prod_{n\in\alpha} M_n \to M$. For each $k \in \omega$, let $\alpha_k \subseteq \alpha$, and let $f_k : \prod_{n\in\alpha_k} M_n \to M$.

We say that f is the **controlled limit** of the sequence $(f_k)_{k\in\omega}$ if for all $r \in \mathbb{R}_+^\alpha$ and $\varepsilon > 0$, there is some finite subset $N_{r,\varepsilon} \subseteq \alpha$ and some index $K_{r,\varepsilon} \in \omega$ such that for all $K \geq K_{r,\varepsilon}$ and for all $x \in \prod_{n\in\alpha} M_n$ with $||x_n|| < r_n$ whenever $n \in N_{r,\varepsilon}$, we have

$$d(f_K(x_n : n \in \alpha_K), f(x)) \leq \varepsilon.$$

Proposition 2.3.4. *Controlled limits of controlled functions are controlled.*

Proof. For each k, let $(\lambda_k, N_k, \delta_k)$ be controllers for f_k. Suppose f is the controlled limit of $(f_k)_{k\in\omega}$.

For all $r \in \mathbb{R}_+^\alpha$ and $\varepsilon > 0$, the controlled limit assumption gives us some $N_{r,\varepsilon}$ and $K_{r,\varepsilon}$ such that in particular, whenever $||x_n|| < r_n$ for $n \in N_{r,\varepsilon}$, we have

$$d(f_{K_{r,\varepsilon}}(x), f(x)) \leq \varepsilon.$$

Define λ by

$$r \mapsto \lambda_{K_{r,1}}(r) + 1.$$

Define N by

$$(r, \varepsilon) \mapsto N_{r,\min(1,\varepsilon/3)} \cup N_{K_{r,\min(1,\varepsilon/3)}}(r, \min(1, \varepsilon/3)).$$

Define δ by

$$(r, \varepsilon) \mapsto \delta_{K_{r,\min(1,\varepsilon/3)}}(r, \min(1, \varepsilon/3)).$$

Let $r \in \mathbb{R}_+^\alpha$ and $\varepsilon > 0$. Without loss of generality, we will suppose $\varepsilon \leq 1$. Suppose $||x_n|| < r_n$ for $n \in N(r, \varepsilon)$. Then we have $||x_n|| < r_n$ for $n \in N_{r,\varepsilon/3} \cup N_{K_{r,\varepsilon/3}}(r, \varepsilon/3)$. So,

$$d(f_{K_{r,\varepsilon/3}}(x), f(x)) \leq \varepsilon/3$$

and

$$||f_{K_{r,\varepsilon/3}}(x)|| \leq \lambda_{K_{r,\varepsilon/3}}(r).$$

By the triangle inequality, we get

$$||f(x)|| \leq \lambda_{K_{r,\varepsilon/3}}(r) + \varepsilon/3 = \lambda(r, \varepsilon).$$

Suppose we also have $||y_n|| < r_n$ and $d(x_n, y_n) < \delta(r, \varepsilon)$ for $n \in N(r, \varepsilon)$. Then we have $||y_n|| < r_n$ and $d(x_n, y_n) < \delta_{K_{r,\varepsilon/3}}(r, \varepsilon/3)$ for $n \in N_{r,\varepsilon/3} \cup N_{K_{r,\varepsilon/3}}(r, \varepsilon/3)$. So we get

$$d(f_{K_{r,\varepsilon/3}}(y), f(y)) \leq \varepsilon/3$$

and

$$d(f_{K_{r,\varepsilon/3}}(x), f_{K_{r,\varepsilon/3}}(y)) \leq \varepsilon/3.$$

By the triangle inequality, we get

$$d(f(x), f(y)) \leq 3(\varepsilon/3) = \varepsilon.$$

□

For countable products, we can check a converse to the previous proposition. This shows how a controlled function on a countable product is a controlled limit of finitary controlled functions.

Proposition 2.3.5. *Let $f : \prod_{n\in\omega} M_n \to M$. Then f is controlled if and only if there is a sequence of controlled functions $f_k : \prod_{n<k} M_n \to M$ such that f is the controlled limit of the sequence $(f_k)_{k\in\omega}$.*

Proof. ($\to$) For each k, define f_k by

$$(x_0, \ldots, x_{k-1}) \mapsto f(x_0, \ldots, x_{k-1}, \star_k, \star_{k+1}, \star_{k+2}, \ldots)$$

where the $\star_i$ are the basepoints of the corresponding spaces. Each f_k is essentially a restriction of f, so it remains controlled.

Let $r \in \mathbb{R}_+^\omega$ and $\varepsilon > 0$. We will refer to controllers λ, N, δ for f. Note that x coincides with the sequence $(x_0, \ldots, x_{K-1}, \star_K, \star_{K+1}, \ldots)$ on coordinates $k < K$. Let $K \geq \max N(r, \varepsilon)$, and suppose x has $||x_n|| < r_n$ for all $n \in N(r, \varepsilon)$. We have that

$$d(f(x^\star), f(x)) \leq \varepsilon$$

where

$$x^\star = (x_0, \ldots, x_{K-1}, \star_K, \star_{K+1}, \ldots)$$

since x^* and x agree on all coordinates in $N(r, \varepsilon)$. But $f(x^\star) = f_K(x)$ by definition, so we have shown that $d(f_K(x), f(x)) \leq \varepsilon$.

This verifies the definition of being a controlled limit using $N_{r,\varepsilon} = N(r, \varepsilon)$ and $K_{r,\varepsilon} = \max N(r, \varepsilon)$.

($\leftarrow$) By the last proposition. □

The next few propositions explain how the definition of controlled functions degenerates to familiar properties in simpler cases involving finite products.

Proposition 2.3.6. *Suppose f is a function whose domain is a finite product of spaces. Then f is controlled iff the restriction of f to any bounded subset is a bounded and uniformly continuous function.*

Proof. Let $f : \prod_{n<K} M_n \to M$.

($\rightarrow$) Let λ, N, δ be controllers for f. Fix a bounded subset B of the domain of f. Choose $r \in \mathbb{R}^K_+$ so that $B \subseteq B_r(\prod M_n)$. Then $x \in B$ implies $||f(x))|| \leq \lambda(r)$ by the definition of λ, and so f is bounded on B. For any $\varepsilon > 0$, we get that $x, y \in B$ and $d(x_n, y_n) < \delta(r, \varepsilon)$ for all $n < K$ implies $d(f(x), f(y)) \leq \varepsilon$ by definition of δ. So f is uniformly continuous on B.

($\leftarrow$) For each $r \in \mathbb{R}^K_+$, let U_r be an upper bound for f on the set $B_r(\prod M_n)$, and let δ_r be a uniform continuity modulus for f on that same set.

Define λ by $r \mapsto U_r$.

Define N to be the constant function returning all K indices of the domain.

Define δ by $(r, \varepsilon) \mapsto \delta_r(\varepsilon)$.

It is trivial to check that f is controlled by these functions. □

Corollary 2.3.7. *Let $f : \prod_{n\in\omega} M_n \to M$. Then f is controlled if and only if f is the controlled limit of some functions $f_k : \prod_{n<k} M_n \to M$, where each f_k restricts to a bounded and uniformly continuous function on any bounded set.*

Proof. The last two propositions. □

Corollary 2.3.8. *Suppose f is a function whose domain is a finite product of bounded spaces. Then f is controlled iff f is bounded and uniformly continuous.*

Corollary 2.3.9. *Suppose f is a function whose domain is a finite product of compact spaces. Then f is controlled iff f is continuous.*

Proof. Compact spaces are bounded, and continuous functions on compact spaces are bounded and uniformly continuous. □

Corollary 2.3.10. *Suppose f is a function whose domain is a finite product of proper spaces. Then f is controlled iff f is continuous.*

Proof. ($\rightarrow$) Since the restriction of f to any bounded set is uniformly continuous, f is continuous.

($\leftarrow$) Closed balls are compact in proper spaces, so the restriction of f to any closed ball is bounded and uniformly continuous. Any bounded subset is contained in some closed ball, so the restrictions of f to bounded subsets are bounded and uniformly continuous. □

Corollary 2.3.11. *Suppose f is a function whose domain is a finite product of spaces with the discrete metric. Then f is controlled iff f is bounded.*

Proof. Every function is uniformly continuous with respect to the discrete metric. □

We end this section with a lemma that appears frequently when proving fundamental results in continuous model theory. The lemma and its argument are slightly adapted from Proposition 2.10 of [3]. It turns an ε-δ relation between two real-valued functions into a continuous, increasing function that lets us bound one of the functions in terms of the other.

Lemma 2.3.12. *Let X be any metric space. Let $f, g : X \to \mathbb{R}$ be functions where for all $\varepsilon > 0$ there is $\delta > 0$ such that for all $x \in X$ we have*

$$f(x) \leq \delta \text{ implies } g(x) \leq \varepsilon.$$

Then there is a continuous, increasing function $\alpha : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ such that

- *$\alpha(0) = 0$, and*
- *for all $x \in X$, we have $g(x) \leq \alpha(f(x))$.*

Proof. First we will obtain a function $\Delta : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ capturing some of the ε-δ behavior. Consider the subset A of $\mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}$ which contains (ε, δ) whenever $f(x) \leq \delta$ implies $g(x) \leq \varepsilon$ for all $x \in X$. We have the following closure properties of A.

- $(\varepsilon, \delta) \in A$ implies $(\varepsilon, \delta') \in A$ whenever $\delta' < \delta$.
- $(\varepsilon, \delta) \in A$ implies $(\varepsilon', \delta) \in A$ whenever $\varepsilon' > \varepsilon$.

If we view $\mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}$ as a quarter-plane with a horizontal axis for ε and a vertical axis for δ, then this means A is closed downward and to the right. That is, A can be viewed as a union of sets of the form $[\varepsilon, \infty) \times [0, \delta]$. The hypotheses ensure that for every $\varepsilon > 0$ there is a positive δ such that A contains a set of this form. It is not hard in this case to construct a continuous, increasing function $\Delta : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ whose graph is inside A and thus satisfies

$$f(x) \leq \Delta(\varepsilon) \text{ implies } g(x) \leq \varepsilon$$

for all $\varepsilon > 0$.

Next, we define a function $\beta : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ by

$$\beta(t) = \inf\{\varepsilon > 0 : \Delta(\varepsilon) > t\}.$$

This β has the following properties:

- $\beta(0) = 0$,
- β is increasing,
- for any $\varepsilon > 0$, whenever $t < \Delta(\varepsilon)$, we have $\beta(t) \leq \varepsilon$,
- $\lim_{t \to 0} \beta(t) = 0$,
- for all $x \in X$, we have $g(x) \leq \beta(f(x))$, since otherwise there would be an x with $\beta(f(x)) < g(x)$, which by definition of β would give an impossible $\varepsilon > 0$ with $\varepsilon < g(x)$ and $\Delta(\varepsilon) > f(x)$.

Finally, we will use β to construct the desired function α. By the above properties of β, it will suffice to make α continuous, increasing, have $\alpha(0) = 0$, and have $\alpha(t) \geq \beta(t)$ for all $t \in \mathbb{R}_{\geq 0}$. Our approach is to construct a piecewise linear function above β. We will start rather arbitrarily at $\alpha(1)$, first defining $\alpha(t)$ for $t \in [0,1]$ and then for $t \in (1, \infty)$.

We start with defining α on $[0,1]$. Let $(t_n : n \in \mathbb{N})$ be a decreasing sequence in $\mathbb{R}_{\geq 0}$ with $t_0 = 1$ and $\lim_{n\to\infty} t_n = 0$. Define the following.

$$\begin{aligned} \alpha(0) &= 0 \\ \alpha(1) &= \beta(2) \\ \alpha(t_n) &= \beta(t_{n-1}) \text{ for } n \geq 1 \end{aligned}$$

Complete the definition of α for $t \in [0,1]$ by making α linear on each segment $[t_{n+1}, t_n]$ with $n \geq 0$. This definition keeps α increasing on $[0,1]$ because β is increasing. To see that $\alpha(t) \geq \beta(t)$ for $t \in [0,1]$, notice that if $t \in [t_{n+1}, t_n]$ for $n \geq 1$, then $t \leq t_n$ and $(t, \alpha(t))$ is on the increasing linear segment from $(t_{n+1}, \beta(t_n))$ to $(t_n, \beta(t_{n-1}))$, so we have $\beta(t) \leq \beta(t_n) \leq \alpha(t)$. For the case where $t \in [t_1, t_0]$, we have that $t \leq 1$ and $(t, \alpha(t))$ is on the segment from $(t_1, \beta(1))$ to $(1, \beta(2))$, and so $\beta(t) \leq \beta(1) \leq \alpha(t)$.

Now we extend α to $(1, \infty)$. Let $(s_n : n \in \mathbb{N})$ be an increasing sequence in $\mathbb{R}_{\geq 0}$ with $s_0 = 1$, $s_1 = 2$, and $\lim_{n\to\infty} s_n = \infty$. We have already defined $\alpha(1) = \beta(2)$, and we continue to define

$$\alpha(s_n) = \beta(s_{n+1}) \text{ for } n \geq 1$$

and then complete the definition of α by making α linear on each segment $[s_n, s_{n+1}]$ with $n \geq 0$. Again, α stays increasing because β is increasing. To see that $\alpha(s) \geq \beta(s)$ for $s \in [1, \infty)$, notice that if $s \in [s_n, s_{n+1}]$ for $n \geq 1$, then $s \leq s_{n+1}$ and $(s, \alpha(s))$ is on the increasing linear segment from $(s_n, \beta(s_{n+1}))$ to $(s_{n+1}, \beta(s_{n+2}))$, so we have $\beta(s) \leq \beta(s_{n+1}) \leq \alpha(s)$. For the case where $s \in [s_0, s_1]$, we have $s \leq 2$ and $(s, \alpha(s))$ is on the segment from $(1, \beta(2))$ to $(2, \beta(s_2))$, and so $\beta(s) \leq \beta(2) \leq \alpha(s)$. $\square$

2.4. **Examples of controlled functions.** The propositions and corollaries from the last section show that many familiar functions are controlled. We call attention to a few in this section.

Example 2.4.1. Any finitary, continuous real function is controlled since $\mathbb{R}$ is proper. In particular, the addition and multiplication maps $+$ and $\cdot$ are controlled.

Example 2.4.2. For any complex Hilbert space H, the addition operator and the map $\lambda_z : H \to H$ defined by scaling by $z \in \mathbb{C}$ are both controlled, since they are bounded and uniformly continuous when restricted to any bounded subset of H^2 or H, respectively.

The next example represents one typical way in which infinitary controlled functions arise.

Example 2.4.3. Consider $f : \mathbb{R}^{\mathbb{N}} \to [0, 2]$ given by defining $f(x_n : n \in \mathbb{N})$ to be

$$\sum_{n\in\mathbb{N}} \frac{\min(1, |x_n|)}{2^n}.$$

This f is controlled. Consider the functions f_N giving the partial sums

$$f_N(x_n : n < N) = \sum_{n<N} \frac{\min(1, |x_n|)}{2^n}.$$

Each f_N is bounded and uniformly continuous, hence controlled. Let $\varepsilon > 0$. Then for all N such that $2^{1-N} < \varepsilon$ and for all x, we have $d(f_N(x), f(x)) \leq \varepsilon$. This verifies that f is the controlled limit of the f_N. Note that we did not need to consider bounds r_n for x_n because of the $\min(1, \cdot)$ compositions.

The next example is a slight variation of the last which does not have as much uniform behavior.

Example 2.4.4. Consider the map $f : \mathbb{R}^{\mathbb{N}} \to \mathbb{R}$ given by defining $f(x_n : n \in \mathbb{N})$ to be

$$x_0 x_1 + \sum_{n\in\mathbb{N}} \frac{\min(1, |x_n|)}{2^n}.$$

This f is controlled. Consider the functions f_N defined by

$$f_N(x_n : n < N) = x_0 x_1 + \sum_{n<N} \frac{\min(1, |x_n|)}{2^n}.$$

Each f_N is controlled since it is bounded and uniformly continuous when restricted to bounded sets. An argument like in the last example shows that f is the controlled limit of the f_N. Notice that we would need bounds on x_0 and x_1 in order to bound f and provide continuity moduli.

The final example shows one way that limits of sequences can appear in the logic we are developing. We will make significant use of functions like this. This construction appears in section 3.2 of [5], under the name of forced limits. The idea is to take an arbitrary sequence (x_n) and inductively project each x_{n+1} into an interval around x_n. This provides a sequence which converges at a known rate. If the original sequence already converged that quickly, the limit will be the same.

Example 2.4.5. Consider the map $f : \mathbb{R}^{\mathbb{N}} \to \mathbb{R}$ given by defining

$$f(x_n : n \in \mathbb{N}) = \lim_{n\to\infty} (x'_n : n \in \mathbb{N})$$

where we define the cauchy sequence x'_n inductively as follows. Let $x'_0 = x_0$. For each $n > 0$, let x'_n be the element in the interval

$$[x'_{n-1} - \frac{|x_0|}{2^n}, x'_{n-1} + \frac{|x_0|}{2^n}]$$

which is closest to x_n.

We now check that f is controlled. For each $N \in \mathbb{N}$, define f_N by

$$f_N(x_0, x_1, \ldots, x_{N-1}) = x'_{N-1}$$

with x'_{N-1} inductively defined as it was above. Each f_N is controlled since it is bounded and uniformly continuous when restricted to bounded subsets. Also, notice that for all x, the value of $f_N(x)$ differs from $f(x)$ by at most

$$\sum_{n \geq N} \frac{|x_0|}{2^n} = |x_0| 2^{1-N}.$$

Let $r_0 \in \mathbb{R}_+$ and $\varepsilon > 0$. For any N such that $|r_0|2^{1-N} < \varepsilon$, whenever we have $|x_0| < r_0$, we get $d(f_N(x), f(x)) \leq \varepsilon$. So f is the controlled limit of the f_N and hence is controlled.

2.5. **Structures and signatures.** We will now define the kinds of objects for which our logic is suited.

Definition 2.5.1. A **metric structure** M is a tuple $(\mathcal{S}_M, \mathcal{F}_M, \mathcal{R}_M)$ such that the following hold:

(1) Each $S \in \mathcal{S}_M$ is a pointed, complete metric space $(M_S, d_S, \star_S)$ where d_S is the metric and $\star_S \in M_S$. We call these spaces in $\mathcal{S}_M$ the **sorts** of the structure M.
(2) Each $f \in \mathcal{F}_M$ is a controlled function from a product of sorts into another sort, i.e.

$$f : \prod_{n \in \alpha} M_{S_n} \to M_S$$

with α any index set. We call these the **functions** of M.
(3) Each $R \in \mathcal{R}_M$ is a controlled function from a finite or countable product of sorts into $\mathbb{R}$, i.e.

$$R : \prod_{n \in \alpha} M_{S_n} \to \mathbb{R}$$

with α any index set. We call these the **relations** of M.

Notation 2.5.2. For convenience, we will often use the same notation for a metric structure M and the underlying metric space of the sort. For example, when there is only one sort M_S, we often just say points in M rather than points in M_S. Similarly, we drop the subscripts and just talk about the metric d and basepoint $\star$ of the structure M.

One goal when applying model theory is to understand a structure by studying related, nicer structures and transporting the information back. To facilitate this, we want a notion of when structures are of the same type so that there is a common language for statements about them. The way we do this is to define an object that keeps track of key features of the structure, such as how many sorts they have and how many functions with what domains, ranges, and controllers. This determines a framework that allows us to ask questions which will be meaningful in any structure fitting that framework.

Definition 2.5.3. A **signature** is a tuple $(\mathcal{S}, \mathcal{F}, \mathcal{R}, \text{domain,range}, \text{control})$ where

(1) $\mathcal{S}$, $\mathcal{F}$, and $\mathcal{R}$ are disjoint sets, whose elements are called the **sort symbols**, **function symbols**, and **relation symbols**, respectively.
(2) domain is a function on $\mathcal{F} \cup \mathcal{R}$ that assigns a sequence of sort symbols.
(3) range is a function on $\mathcal{F}$ that assigns a sort symbol.
(4) control is a function on $\mathcal{F} \cup \mathcal{R}$ that assigns controllers (λ, N, δ).

Informally, we want to say that a structure fits a signature if there is a way to pair each sort, function, and relation with a corresponding symbol in a way that respects the domains, ranges, and controllers.

Definition 2.5.4. Given a signature $L = (\mathcal{S}, \mathcal{F}, \mathcal{R}, \dots)$, an L-**structure** is a metric structure $M = (\mathcal{S}_M, \mathcal{F}_M, \mathcal{R}_M)$ such that there are maps $\mathcal{S} \to \mathcal{S}_M$, $\mathcal{F} \to \mathcal{F}_M$, and $\mathcal{R} \to \mathcal{R}_M$ satisfying the conditions below.

First, we establish some notation. For each function symbol $f \in \mathcal{F}$ and relation symbol $R \in \mathcal{R}$, the maps above determine a corresponding function and relation in the structure M, which we denote by f^M and R^M respectively. For each sort symbol S, we will denote the corresponding sort by M_S.

For each $f \in \mathcal{F}$ we must have

- domain$(f) = (S_n : n \in \alpha)$ iff the domain of f^M is $\prod_{n\in\alpha} M_{S_n}$,
- range$(f) = S$ iff the range of f^M is M_S,
- control(f) are controllers for f^M,

and similarly for each $R \in \mathcal{R}$ we must have

- domain$(R) = (S_n : n \in \alpha)$ iff the domain of R^M is $\prod_{n\in\alpha} M_{S_n}$,
- control(R) are controllers for R^M.

The benefit of this abstraction is that this lets us write down formal expressions made of symbols from L. These can be assigned a meaning in any L-structure. This is analogous to treating polynomials as formal objects and varying the field over which we consider them. For example, we can consider $x \cdot x + 1$ over $\mathbb{R}$, $\mathbb{C}$, or $\mathbb{Z}_2$. We can ask questions like whether it has a root, comparing the answers in the different number systems. The abstraction to structures and signatures gives us a formal way to work like this.

2.6. **Languages and formulas.** In this section we inductively define a formal language for each signature.

If the reader is unfamiliar with presentations of logic, keep in mind that the symbols in a signature are formal objects and are not actually functions. We will be discussing sequences of symbols at a level detached from any particular metric structure. The notation is suggestively chosen so that these sequences look like usual functions on a structure, because later we will specify a structure and assign an actual function. This is similar to how the symbol $+$ and expression $x + y$ do not meaningfully correspond to functions until we decide where we are working. It could be real addition, or matrix addition, or an arbitrary binary function which has nothing to do with any usual kind of addition.

If the reader is familiar with discrete logic, the transition to a more analytic setting makes for some initially strange-looking changes. In place of the familiar boolean connectives "and", "or", "implies", and so on, we have compositions with a seemingly large family of real-valued functions. In place of the familiar $\forall$ and $\exists$ quantifiers, we have decorated sup and inf quantifiers. We will see in the next section how these are interpreted in a structure.

We keep careful track of sort symbols in the following definition. This is in order to avoid oddities like composing symbols when their domains and ranges do not coincide. Each object introduced below is assumed to be distinct from the others.

Definition 2.6.1. Given a signature $L = (\mathcal{S}, \mathcal{F}, \mathcal{R}, ...)$, its **language**, which we continue to denote L, consists of the following.

The L-**variables** consist of infinitely many objects $v_1, v_2, v_3, \ldots$ for each sort symbol S. If x is a variable for sort S then we say x is of sort S.

The L-**basepoint symbols** consist of a symbol $\star_S$ for each sort S.

The L-**distance symbols** consist of a symbol d_S for each sort S.

The L-**terms** are inductively defined as follows.

(1) For each sort symbol S, each variable

$$x$$

of sort S is a term of sort S

(2) For each sort symbol S, the basepoint symbol

$$\star_S$$

is a term of sort S.

(3) If t_n for $n \in \alpha$ are terms of sort S_n respectively, and $f \in \mathcal{F}$ has $\mathrm{domain}(f) = (S_n : n \in \alpha)$ then

$$f(t_n : n \in \alpha)$$

is a term of sort $\mathrm{range}(f)$.

The **atomic** L-**formulas** are defined as follows.

(1) If t_1, t_2 are both terms of the same sort S, then

$$d_S(t_1, t_2)$$

is an atomic L-formula.

(2) If t_n for $n \in \alpha$ are terms of sort S_n respectively, and R is a relation symbol with $\mathrm{domain}(R) = (S_n : n \in \alpha)$, then

$$R(t_n : n \in \alpha)$$

is an atomic L-formula.

The L-**formulas** are inductively defined as follows. We simultaneously define the **free variables** of a formula.

(1) Each atomic L-formula ϕ is an L-formula, and all variables appearing in ϕ are free in ϕ.
(2) If ϕ_n for $n \in \alpha$ are L-formulas and u is a controlled real-valued function $u : \mathbb{R}^\alpha \to \mathbb{R}$, then
$$u(\phi_n : n \in \alpha)$$
is an L-formula. If a variable x is free in some ϕ_n, then it is free in $u(\phi_n : n \in \alpha)$ as well.
(3) If ϕ is an L-formula, x is a finite tuple of variables, and r, r' are both real tuples of the same length as x, then
$$\sup_x]_r^{r'} \phi$$
and
$$\inf_x]_r^{r'} \phi$$
are L-formulas. The variables in x are not free in these formulas. If y is a variable which is free in ϕ and not a member of the tuple x, then y is also free in these formulas.

If ϕ is constructed without any application of step (3), that is, if no sup or inf appears in ϕ, we say ϕ is **quantifier-free**.

If ϕ is an L-formula with no free variables, then we say that ϕ is an L-**sentence**.

Definition 2.6.2. If L and L' are languages such that every sort and symbol from the signature of L is also in L' with the same control, domain, and range values, then we say L' is an **extension** of L, and we write $L \subseteq L'$. We also say in this case that L is a **reduct** of L'.

With L,L' as above, if M' is an L'-structure and M is an L-structure obtained from M' by ignoring the sorts and symbols outside of L, then we call M' an **extension** of M to L', and we call M a **reduct** of M' to L.

2.7. **Interpretation of terms and formulas.** Given an L-structure M, we explain in this section how to assign to every L-term or L-formula ϕ a function $\phi^M : \prod_{n\in\alpha} M_n \to \mathbb{R}$. We will see that such ϕ^M are controlled functions, and moreover that controllers (λ, N, δ) can be determined inductively just from the signature of L. Hence the controllers obtained actually control ϕ^M for every choice of M.

First we define interpretation of terms.

Definition 2.7.1. Given an L-structure M and an L-term t, we inductively define the **interpretation** t^M of t in M as follows.

(1) If x is a variable of sort S, then x^M is the identity function $M_S \to M_S$.
(2) If $\star_S$ is the basepoint symbol for S, then $\star_S^M$ is the constant 0-ary function $M_S^0 \to M_S$ which maps to the basepoint of M_S.
(3) If $f(t_n : n \in \alpha)$ is a term where f is a function symbol and each t_n^M is already defined, then $(f(t_n : n \in \alpha))^M$ is the function $f^M(t_n^M : n \in \alpha)$.

Now we extend the definition of interpretation to formulas.

Definition 2.7.2. Given an L-structure M and an L-formula ϕ, we inductively define the **interpretation** ϕ^M of ϕ in M as follows.

(1) If ϕ is an atomic formula $d_S(t_1, t_2)$ where t_1^M and t_2^M are already defined and d_S is the distance symbol for the sort S, then ϕ^M is $d_S^M(t_1^M, t_2^M)$.
(2) If ϕ is an atomic formula $R(t_n : n \in \alpha)$ where each t_n^M is already defined, then ϕ^M is the real-valued function $R^M(t_n^M : n \in \alpha)$.
(3) If ϕ is of the form $u(\psi_n : n \in \alpha)$ where each ψ_n^M is already defined, then ϕ^M is the real valued function $u(\psi_n^M : n \in \alpha)$.
(4) Suppose ϕ is of the form $\sup_x]_r^{r'}\psi$, with finite variable tuple $x = (x_k)_{k<K}$ having $x_0, \ldots, x_{K-1}$ of sorts $S_0, \ldots, S_{K-1}$ respectively, real tuples $r = (r_k)_{k<K}$ and $r' = (r'_k)_{k<K}$, and where ψ^M is already defined. If $r_k \neq r'_k$ for all $k < K$, we define ϕ^M to be the real-valued function

$$\left(\prod_{k<K} \frac{1}{|r'_k - r_k|}\right) \cdot \int_{r_{K-1}}^{r'_{K-1}} \cdots \int_{r_1}^{r'_1} \sup_{x \in B_{(\rho_k : k<K)}(M_S)} \psi^M(x) d\rho_1 \cdots d\rho_{K-1}$$

where $M_S = \prod_{k<K} M_{S_k}$. If any $r_k = r'_k$, we just define ϕ^M to be the constant 0 function. This makes ϕ^M the average supremum as we vary the radii of the balls between bounds given in r and r'. We verify in propositions below that this is a Riemann integral with some nice properties determined by the language. We also check that it is equivalent to use either open or closed balls above.
(5) If ϕ is of the form $\inf_x]_r^{r'}\psi$, we handle it as in the $\sup_x]_r^{r'}\psi$ case, but with inf in place of each sup.

Note that the quantifier cases give functions on a domain which is typically a product with fewer factors factor than the domain of ψ^M. For example, if x was the only free variable in ψ, so that ϕ no longer has free variables, then $\sup_x]_r^{r'}\psi$ interprets as a constant 0-ary function $M_S^0 \to \mathbb{R}$.

As promised, we check the claims made at the start of this section and in the definition.

Lemma 2.7.3. ***(Terms are controlled)*** *For every L-term t, there are controllers (λ, N, δ) such that for all L-structures M, the interpretation t^M is controlled by (λ, N, δ).*

Proof. The proof is by induction. Variables interpret as identity functions, and basepoint symbols interpret as constant functions with 0 magnitude, so both are easily checked. Since compositions of controlled functions are controlled, any term of the form $f(t_k : k \in \alpha)$ is controlled. The construction of controllers for the composition can be done with controllers given by L, and so is independent of the structure over which we work. □

Lemma 2.7.4. ***(Formulas are controlled)*** *For every L-formula ϕ, there are controllers (λ, N, δ) such that for all L-structures M, the interpretation ϕ^M is controlled by (λ, N, δ).*

Proof. We continue with induction on formulas, having established the case of terms above. Nothing is substantially different for the cases of atomic formulas and composition with controlled $u : \mathbb{R}^\alpha \to \mathbb{R}$. So, we will only discuss the quantifier case. Since inf is similar to sup, we will only discuss sup. For convenience, we will moreover assume the quantification is over a single variable. The finite tuple case is an easy generalization, but would require an extra set of indices.

Suppose ϕ is of the form $\sup_x]_{R_1}^{R_2}\psi$ where $\psi^M : \prod_{n\in\alpha} M_n \to M_\psi$ is controlled by $(\lambda_\psi, N_\psi, \delta_\psi)$. The case $R_1 = R_2$ is trivial, and the case $R_1 > R_2$ follows from the $R_1 < R_2$ case by composition with multiplication by -1. So we assume $R_1 < R_2$. If x is not free in ψ, then ϕ^M is just ψ^M and we are done. So we assume x is free in ψ. Let S be the sort of x, and for convenience assume the indexing is such that $M_0 = M_S$. This lets us view ϕ as a function $\prod_{n\in\alpha_+} M_n \to M_\psi$, where α_+ is $\alpha - \{0\}$. Let $y = (y_n : n \in \alpha_+)$ be the sequence of remaining free variables in ϕ.

For any $r = (r_n : n \in \alpha_+) \in \mathbb{R}_+^{\alpha_+}$, define $r^* = (r^*_n : n \in \alpha)$ as follows.

$$r^*_n = \begin{cases} R_2 + 1 & \text{when } n = 0 \\ r_n & \text{otherwise} \end{cases}$$

Define λ by

$$\lambda(r) = \lambda_\psi(r^*).$$

Define N by

$$N(r, \varepsilon) = N_\psi(r^*, \varepsilon) - \{0\}.$$

Define δ by

$$\delta(r, \varepsilon) = \delta_\psi(r^*, \varepsilon).$$

We check that these control ϕ^M. Let $r = (r_n : n \in \alpha_+) \in \mathbb{R}_+^{\alpha_+}$ and $\varepsilon > 0$.

Suppose $||y_n|| < r_n$ for all $n \in N(r, \varepsilon)$. Then whenever $||x|| \leq R_2$ (hence $||x|| < r^*_0$) we get $||\psi(x, y)|| \leq \lambda_\psi(r^*) = \lambda(r)$. This implies

$$\sup_{x\in B_\rho(M_0)} \psi^M(x, y) \leq \lambda(r)$$

for all $\rho \leq R_2$, and so

$$\frac{1}{R_2 - R_1} \int_{R_1}^{R_2} \sup_{x\in B_\rho(M_0)} \psi(x, y) d\rho \leq \lambda(r)$$

as required.

Suppose $||y_n|| < r_n$, $||z_n|| < r_n$, and $d(y_n, z_n) < \delta(r, \varepsilon)$ for all $n \in N(r, \varepsilon)$. Then whenever $||x|| \leq R_2$, we get $d(\psi(x, y), \psi(x, z)) \leq \varepsilon$. So for all $\rho \leq R_2$, we have that $\sup_{x\in B_\rho(M_0)} \psi^M(x,y)$ is within ε of $\sup_{x\in B_\rho(M_0)} \psi^M(x,z)$. Thus,

$$\begin{aligned}
& \left|\phi^M(y) - \phi^M(z)\right| \\
= & \frac{1}{R_2 - R_1}\left|\int_{R_1}^{R_2}\left(\sup_{x\in B_\rho(M_0)} \psi^M(x,y) - \sup_{x\in B_\rho(M_0)} \psi^M(x,z)\right) d\rho\right| \\
\leq & \frac{1}{R_2 - R_1}\int_{R_1}^{R_2}\left|\sup_{x\in B_\rho(M_0)} \psi^M(x,y) - \sup_{x\in B_\rho(M_0)} \psi^M(x,z)\right| d\rho \\
\leq & \frac{1}{R_2 - R_1}\int_{R_1}^{R_2} \varepsilon d\rho \\
= & \varepsilon
\end{aligned}$$

as required.

This verifies that ϕ^M is controlled, and we constructed (λ, N, δ) independently of M. □

We will extend our terminology a bit to more conveniently refer to controllers that work independently of the L-structure.

Definition 2.7.5. Given an L-formula ϕ, we say ϕ is **controlled** by (λ, N, δ) if for every L-structure M, every ϕ^M is controlled by (λ, N, δ).

So what we have proved is that every formula is controlled in this sense.

Now we will check that the integrals in the interpretation of quantified formulas are well-defined. We also show a useful result about how L determines partition data and rates of convergence for Riemann sums approximating these integrals.

Lemma 2.7.6. ***Riemann Integrals and Good Partitions***
Let α be a set, let x be a variable of sort M_0, and let $\alpha_+ = \alpha - \{0\}$. Let $\psi(x, y_n : n \in \alpha_+)$ be an L-formula, and let $r < r'$ be reals.

Then we have the following.

- *Let M be any L-structure, and let $b = (b_n : n \in \alpha_+)$ with each b_n in the sort of M corresponding to y_n.*
  - *The function $s_b^M : \mathbb{R}^{\geq 0} \to \mathbb{R}$ defined by*
  $$s_b^M(\rho) = \sup_{x\in B_\rho(M_0)} \psi^M(x,b)$$
  is Riemann integrable.
  - *The function $\bar{s}_b^M : \mathbb{R}^{\geq 0} \to \mathbb{R}$ defined by*
  $$\bar{s}_b^M(\rho) = \sup_{x\in \bar{B}_\rho(M_0)} \psi^M(x,b)$$
  is Riemann integrable.
  - *The integrals $\int_r^{r'} s_b^M(\rho)d\rho$ and $\int_r^{r'} \bar{s}_b^M(\rho)d\rho$ are equal.*

- *For any $R = (R_n : n \in \alpha_+)$ and $\varepsilon > 0$, there are $N \subseteq \alpha_+$, $\Delta > 0$, and a partition $\rho_0 < \cdots < \rho_K$ of $[r, r']$ with $|\rho_k - \rho_{k+1}| = \Delta$ for all $k < K$, such that for any L-structure M, any $b = (b_n : n \in \alpha_+)$ with $||b_n|| < R_n$ for all $n \in N$, and any $s_0^*, \ldots, s_{K-1}^* \in \mathbb{R}$ satisfying*
$$s_b^M(\rho_k) \le s_k^* \le s_b^M(\rho_{k+1})$$
for all $k < K$, we have
$$\left| \int_r^{r'} s_b^M(\rho) d\rho - \sum_{k<K} s_k^* \Delta \right| < \varepsilon.$$
*We call this partition an (L, R, ε)-**good partition** for the formula $\sup_x]_r^{r'} \psi$.*

Similar statements hold when sup *is replaced by* inf*, and when x is a finite tuple of variables.*

Proof. Let (λ, N, δ) control ψ. Let R be as in (2). Define R^* to be the sequence given by $R_0^* = r' + 1$ and $R_n^* = R_n$ otherwise. Then provided $||x|| \le r'$ and $||y_n|| < R_n$ for all $n \in N(R, \varepsilon) - \{0\}$, we have $||\psi^M(x, y)|| \le \lambda(R^*)$. Thus, whenever $||y_n|| < R_n$ for all $n \in N(R, \varepsilon) - \{0\}$, the function $\sup_{x \in B_\rho(M_S)} \psi^M(x, y)$ is bounded by $\lambda(R^*)$ and monotonic on $[r_1, r_2]$. This is sufficient to carry out the standard development of the Riemann integral. Since the bounds are determined by L, the partitions involved in the Riemann sums approximating the integrals can be chosen independently of the L-structure M. This verifies $(1a)$, $(1b)$, and (2), so we just need to check $(1c)$.

Notice that for all $\rho \in [r, r']$ and $\varepsilon > 0$, we have
$$\sup_{x \in B_\rho} \psi \le \sup_{x \in \bar{B}_\rho} \psi \le \sup_{x \in B_{\rho+\varepsilon}} \psi \le \sup_{x \in \bar{B}_{\rho+\varepsilon}} \psi.$$
Integrating each term gives us
$$\int_r^{r'} \sup_{x \in B_\rho} \psi d\rho \le \int_r^{r'} \sup_{x \in \bar{B}_\rho} \psi d\rho \le \int_r^{r'} \sup_{x \in B_{\rho+\varepsilon}} \psi d\rho \le \int_r^{r'} \sup_{x \in \bar{B}_{\rho+\varepsilon}} \psi d\rho.$$
But by changing variables this is the same as
$$\int_r^{r'} \sup_{x \in B_\rho} \psi d\rho \le \int_r^{r'} \sup_{x \in \bar{B}_\rho} \psi d\rho \le \int_{r+\varepsilon}^{r'+\varepsilon} \sup_{x \in B_\rho} \psi d\rho \le \int_{r+\varepsilon}^{r'+\varepsilon} \sup_{x \in \bar{B}_\rho} \psi d\rho.$$
This holds for all $\varepsilon > 0$, and the integrals are continuous with respect to their endpoints, so we get the claim. □

2.8. **Some comments on the quantifiers.** Since our langauge includes composition with multiplication by any real number, we could just as well have chosen to forego the scaling by $\frac{1}{r_2 - r_1}$ when defining the quantifiers. The average seems more natural when explaining the semantics.

An important observation about the quantifiers is that given an L-formula ψ and any $r, \varepsilon > 0$, we will know that the interpretation of $\sup_x]_r^{r+\varepsilon} \psi$ in any

M is between $\sup_{x\in\bar{B}_r(M)} \psi^M$ and $\sup_{x\in\bar{B}_{r+\varepsilon}(M)} \psi^M$. At first glance, this might seem like enough to approximate the actual value of $\sup_{x\in\bar{B}_r(M)} \psi^M$ arbitrarily well, but issues arise when $\sup_{x\in\bar{B}_\rho(M)} \psi^M(x)$ is discontinuous as a function of ρ around $\rho = r$. This potential discontinuity makes us unable to generally guarantee the existence of an L-formula whose interpretation is within ε of $\sup_{x\in\bar{B}_r(M)} \psi^M$ across all L-structures. The discontinuity of this sup function is the motivation for introducing an integral into the quantifiers.

However, in many natural classes of structures, each $\sup_{x\in\bar{B}_\rho(M)} \psi^M$ is continuous with respect to ρ, and moreover this continuity can often be described independently of the structure in the class. For example, Hilbert spaces and geodesic spaces have this property. In such classes, we can obtain a formula which interprets as the function $\sup_{x\in\bar{B}_r(M)} \psi^M$ for each M in the class. We can obtain it using a forced limit connective u composed with formulas $\sup_x]_r^{r+\delta}\psi$ when this continuity is understood well enough. We will see another way to obtain this function in nice classes when we discuss definable sets. Readers familiar with bounded continuous logic will see how this relates to the use of sup quantifiers in that setting.

2.9. **Multivariable quantifiers.** In discrete logic, we often quantify over multiple variables at once and do not need to give much thought to the distinction between $\forall x \in M, \forall y \in N$ and $\forall (x,y) \in M \times N$. Similarly, in bounded continuous logic we just have sup and inf quantifiers without the averaging, and we can use that $\sup_x \sup_y$ is the same as $\sup_{(x,y)}$. However, in the current setting we need to be a bit more careful. In general,

$$\operatorname*{avg}_{\rho\in[r,r']} \sup_{x\in B_\rho(M)} \left(\operatorname*{avg}_{\tau\in[t,t']} \sup_{y\in B_\tau(M)} \psi^M(x,y) \right)$$

is not the same as

$$\operatorname*{avg}_{(\rho,\tau)\in[r,r']\times[t,t']} \sup_{(x,y)\in B_\rho(M)\times B_\tau(M)} \psi^M(x,y).$$

For a simple example, consider that in general

$$\max_{x\in\{0,1\}} (f(x,0)+f(x,1)) \neq \max_{x\in\{0,1\}} f(x,0) + \max_{x\in\{0,1\}} f(x,1).$$

This is demonstrated by the function $f(x,y)$ defined by the following.

$$\begin{array}{rcl} (0,0) & \mapsto & 0 \\ (0,1) & \mapsto & 0 \\ (1,0) & \mapsto & -1 \\ (1,1) & \mapsto & 1 \end{array}$$

The language we have defined provides us naturally with both of the above kinds of quantification. The first arises from interpreting nested quantifiers as in the formula

$$\sup_x]_r^{r'} \sup_y]_t^{t'} \psi(x,y).$$

The other cases such as

$$\sup_{x,y}]_{(r,t)}^{(r',t')} \psi(x,y)$$

are explicitly provided as a single quantification step in the definition of the language.

These tuple quantifiers can be expressed using the nested cases by partitioning the box determined by (r,t) and (r',t') and building the required multivariable integral as a limit of Riemann sums. However, the rate of convergence of these Riemann sums will depend on the remaining free variables. Because of this, attempting to express this limit via composition, for example by composing with a forced limit function u, will only result in a formula guaranteed to be accurate given bounds on the free variables.

2.10. **Quantifying over sequences.** Since we can compose with infinitary real functions in our formulas, our logic has formulas which quantify over infinitely many variables in a certain way. Of course, we cannot quantify in an arbitrary way over sequences. We are limited to what can be built using our finitary quantifiers and composition with (controlled) connectives. But this still includes some important limited versions of infinite quantification which will play an important role in our applications.

For example, the following is a valid sentence σ in any language.

$$\sum_{n\in\mathbb{N}} \frac{\sup_{x_n}]_0^1 d(x_n,\star)}{2^n}$$

This sentence contains countably many variables, and none of them are free. It is easy to come up with other examples like this by taking weighted sums of finitely quantified formulas.

We can think of the above σ as a sum of sentences about single variables. Notice that $\sigma^M = 0$ iff for all n, the ball $B_1(M_{S_n})$ contains only its basepoint $\star_{S_n}^M$. But we can also view σ as a sentence about countable sequences. That is, it is also true that $\sigma^M = 0$ iff the only sequence in $\prod B_1(M_{S_n})$ is $(\star_{S_n}^M : n \in \mathbb{N})$.

Informally, we can get away with this infinite quantification and still get expected results in model theory because given any $\varepsilon > 0$, the value of the above sum can be determined to within ε using only finitely many variables. In other words, the quantification can be thought of as finite up to ε.

Next, we will look at another approach to infinite quantification formulas by using forced limits and knowledge of controllers. Again, this is possible because being controlled essentially means that a formula ψ only depends up to ε on some known finite subset of its coordinates.

To clarify the meaning of the limit taken in the following proposition, recall that any interpretation ψ^M of a formula ψ with free variables indexed by α

is some function $M^\alpha \to \mathbb{R}$. When we quantify over some finite tuple x of variables indexed by $\{1, \dots, N\}$, we get a formula ψ_N given by

$$\sup_x]_r^{r'}\psi$$

whose interpretation $(\psi_N)^M$ is a function $M^{\alpha - \{1,\dots,N\}} \to \mathbb{R}$. This function $(\psi_N)^M$ extends uniquely to a function $M^\alpha \to \mathbb{R}$ which is constant with respect to the coordinates $\{1, \dots, N\}$, defined by setting the value on $m \in M^\alpha$ to be the value of $(m_n : n \in \alpha - \{1, \dots, N\})$ under $(\psi_N)^M$. So, for the purposes of the following definition, we can consider the lim to be a limit of functions all viewed as having a common domain.

Proposition 2.10.1. *Suppose we have an L-formula ψ, countable sequence of variables $(x_n : n \in \alpha)$ of sorts S_n respectively, reals $(R_k : k \in \beta)$ corresponding to each remaining free variable y_k of ψ, and countable real sequences (r_n) and (r'_n) with $r_n < r'_n$ for all n.*

There is an L-formula ϕ such that in every L-structure M, the interpretation ϕ^M satisfies

$$\phi^M = \lim_{N\to\infty}\left(\sup_{(x_1,\dots,x_N)}]_{(r_1,\dots,r_N)}^{(r'_1,\dots,r'_N)}\psi\right)^M$$

when restricted to a subset where $||y_k|| \leq R_k$ for all k.

Proof. We will write $r \cup R$ for the indexed set of reals $\{r'_n : n \in \alpha\} \cup \{R_k : k \in \beta\}$.

Let (λ, N, δ) be the controllers for ψ, and for each $\varepsilon > 0$, let $\psi_{N(r' \cup R, \varepsilon)}$ denote the formula $\sup_{x^*}]_{r^*}^{r'^*}\psi^*$, where

- ψ^* is obtained from ψ by replacing all variables having coordinates outside $N(r' \cup R, \varepsilon)$ with the basepoint of their respective sort,
- x^* is the tuple of variables x_n whose indices appear in $N(r' \cup R, \varepsilon)$, and
- r^* and r'^* are the corresponding tuples for the variables in x^* from among r_n and r'_n.

By definition of the controllers, the value of ψ^M can be determined to within ε by just the coordinates in $N(r' \cup R, \varepsilon)$, provided they satisfy the corresponding bounds in $r' \cup R$. This implies convergence of $\psi^M_{N(r'\cup R,\varepsilon)}$ as $\varepsilon \to 0$, but even stronger, for all $\varepsilon > 0$ and L-structures M, there is some n_ε such that all $\psi^M_{N(r'\cup R, n^{-1})}$ with $n \geq n_\varepsilon$ are within ε of the limit function in M. Since this rate of convergence is determined just from L, we can use a controlled $u : \mathbb{R}^\mathbb{N} \to \mathbb{R}$ to compute the limit. The composition $u(\psi_{N(r'\cup R, n^{-1})} : n \geq 1)$ gives us the L-formula claimed in the proposition. $\square$

For example, this last proposition gives an unambiguous formula in every L-structure equivalent to an expression like

$$\sup_x]_r^{r'}\sum_{n\in\mathbb{N}}\frac{\max(1, d(x_n, \star))}{2^n}$$

where x, r, and r' are countable sequences.

There are clear ways to generalize the above proposition to cases with uncountably many variables, since we can similarly build an increasing collection of finite subsets of coordinates, but we would need to pick a more general notion of limit.

2.11. **Logical distance between formulas.** In discrete logic there is a notion of logical equivalence of formulas. This is typically defined either proof theoretically, by saying that $\phi \leftrightarrow \psi$ is provable with no additional assumptions, or defined semantically by saying that for every L-structure, ϕ^M and ψ^M have equivalent interpretations. In this section we define the analogous notion for formulas in the continuous setting.

Definition 2.11.1. Let $\phi(x)$ and $\psi(x)$ be L-formulas, where x is a finite or contable sequence of variables. We define the **logical distance** between ϕ and ψ to be the supremum of the values $|\phi^M(a) - \psi^M(a)|$ across L-structures M and tuples a from M.

Since our formulas are potentially unbounded, the logical distance between ϕ and ψ could be ∞. For example, if ϕ is just the 0 formula, and ψ is $d(\star, x)$. So this notion of distance is extended real valued. Moreover, two formulas may have a logical distance of 0. A trivial example is the pair of formulas $d(x, \star)$ and $\max(0, d(x, \star))$. These are technically distinct formulas, but their logical distance is easily seen to be 0.

Definition 2.11.2. If ϕ and ψ have a logical distance of 0, we say that they are **logically equivalent**.

At best, we could describe this logical distance as an extended real valued pseudometric. Symmetry and the triangle inequality can be easily verified.

The other interesting behavior to discuss is that we can find pairs of formulas which are arbitrarily close with respect to logical distance. A trivial example of this is that for any formula $\phi(x)$, and any $\varepsilon > 0$, there is also a formula $\phi(x) + \varepsilon$. These formulas will have a logical distance of ε.

Moreover, we can look for dense sets of formulas. This is important to notice because the notion of the size of a language L in the continuous setting is a bit subtle. We will see when we get to some of the basic theorems from model theory that we are not so much interested in the cardinality of the set of L-formulas as we are in the cardinality of a dense subset.

Definition 2.11.3. We say a set F of L-formulas is **dense** if for every L-formula ϕ and every $\varepsilon > 0$, there is some ϕ_ε in F such that the logical distance between ϕ and ϕ_ε is $\leq \varepsilon$.

Notably, while the cardinality of the set of L-formulas will be uncountable for any L, for example because we allow composition with any continuous $u : \mathbb{R} \to \mathbb{R}$, there will often be a countable, dense subset of formulas. Since the cardinality of these dense subsets will be the more important notion, we make the following definition.

Definition 2.11.4. We define the **density** of a language L to be the smallest cardinality among the dense subsets of L-formulas. We denote the density by $\mathrm{density}(L)$.

2.12. **Bounded or discrete structures.** In this section, we will make some observations about our logic applied to structures with bounded sorts. Notably, we will see how the $\sup]_r^{r'}$ and $\inf]_r^{r'}$ quantifiers degenerate to ordinary sup and inf quantifiers over the entire bounded space. Then we will see how bounded continuous logic and discrete logic can be viewed as special cases of our logic.

Proposition 2.12.1. *Let $r \in [D, \infty)$, and let ϕ be an L-formula of the form $\sup]_r^{r+1}\psi$. Then in any L-structure M where the sort M_S is bounded with diameter D, the interpretation ϕ^M is the function $\sup_{x \in M_S} \psi^M$. A similar statement holds with* sup *replaced by* inf.

Proof. By definition, the interpretation ϕ^M is $\int_r^{r+1} \sup_{x \in B_\rho(M_S)} \psi^M dr$. This is the average supremum of $\psi^M(x)$ over the balls of radius $\rho \in [r, r+1]$. But each $\rho \in [r, r+1]$ is at least D, so $B_\rho(M_S)$ is all of M_S. This means $\sup_{x \in B_\rho(M_S)} \psi^M(x)$ is constantly equal to $\sup_{x \in M_S} \psi^M(x)$ for these ρ, and its average value is just $\sup_{x \in M_S} \psi^M(x)$. □

The comments in the rest of this section assume familiarity with bounded continuous logic and discrete logic. It is not completely straightforward that bounded continuous logic and discrete logic are special cases of the current logic, but a few comments make it clear in what sense this is true.

First, note that any discrete structure can be viewed as a discrete metric space by giving it the metric d defined by $d(x, y) = 0$ when $x = y$ and $d(x, y) = 1$ otherwise. This d essentially plays the role of equality. Any relation can be viewed as a function to $\{0, 1\}$ rather than $\{\mathrm{True}, \mathrm{False}\}$. In light of the previous proposition, we can quantify with sup and inf over the whole structure. Since the metric is discrete, sup and inf values will have to be actually realized. This makes any sup or inf equivalent to a discrete formula using $\forall$ or $\exists$. In the context of a discrete metric, our connectives include analogs of finite boolean connectives. In discrete spaces, the interpretations of unbounded continuous logic formulas can degenerate and account for all of the expected discrete formulas.

Continuous logic has some extra formulas however. For example, consider a discrete structure M with a subset A consisting of the realizations of a countable type p. In discrete logic, A might not be definable by a single formula. However, in continuous logic, if A is $\{x \in M : \phi_n(x) = 0 \text{ for all } n \in \mathbb{N}\}$, then we can also realize A as the zero set of the single formula $\sum_{n \in \mathbb{N}} \phi_n(x) \cdot 2^{-n}$.

Our setting also requires fixing a basepoint for each sort. A bounded or discrete structure might not have a natural constant to serve as this basepoint.

One way to address this is to amalgamate the logics, forming a logic with two kinds of sorts: bounded with no basepoint but an assigned diameter, and unbounded with a basepoint. Another ad hoc approach is the following way. Say (M, d) is a metric space with diameter D. Define (M', d') by adding a new point to get $M' = M \cup \{\star\}$ and extending d by setting $d(x, \star) = D + 1$ for all $x \in M$. When we study definable sets, it will be clear that $M = M' - \{\star\}$ is definable in M' and hence can be quantified over as though it were a bounded sort in M'. This observation allows us to translate between bounded or discrete formulas about M and unbounded formulas about M'.

2.13. **Ben Yaacov's unbounded continuous logic.** The quantifiers $\sup_x]_r^{r'}$ and $\inf_x]_r^{r'}$ in our setting are interpreted as an averaged supremum or infimum over balls with radii varying in some bounded range. The motivation comes from wanting our logic to behave well with respect to the ultraproduct construction, i.e. satisfy a version of Łoś's Theorem, which we prove as Theorem 3.7.4. The issue being address is that in general the function $\sup_{x \in B_r(M)} \phi(x)$ is not continuous with respect to r. We get around this by integrating to smooth things out.

The interested reader can contrast this approach with the approach in [4]. There, the author achieves the same goal (i.e. a logic for unbounded metric spaces with this expressivity) in a different manner. He uses sup and inf quantifiers with the variable ranging over the entire sort, and without any integration. Doing so requires restricting the formulas to which the sup and inf are allowed to apply. The author develops syntactic notions of eventual constancy and boundedness of formulas, and additionally modifies the semantics for sup and inf so that they account for an ideal point ∞ at which a formula takes on its eventual constant value. He later develops formulas which interpret as an approximate version of quantification over balls, noting its convenience for stating axioms. For example, these are formulas $\sup_x^{r,r'} \phi$ with the property that

$$\sup_{\nu(x) \leq r} \phi \leq \sup_x^{r,r'} \phi \leq \sup_{\nu(x) < r'} \phi$$

when evaluated in any structure, where $\nu(x)$ is the gauge of x, a more general notion of magnitude which is sufficient for discussing continuity and boundedness restrictions needed for these logics.

Our quantifiers $\sup_x]_{r_1}^{r_2}$ could be constructed in that setting using uniform limits of formulas involving the approximate quantifiers $\sup_x^{r,r'}$ mentioned above. One would simply use Riemann sums, paying attention to the rate of convergence ensured by the properties demanded of the formula ϕ. Conversely, it is not difficult to see:

(1) how our $\sup_x]_{r_1}^{r_2}$ quantifiers could be used to obtain "space-wide" quantification over the formulas with the eventual constancy property, or to

(2) obtain formulas with the same property as approximate quantifiers above.

For (1), average the supremum of an eventually constant formula over sufficiently large r_1, r_2. For (2), average over radii between r_1 and r_2 to obtain the desired inequalities. There are various assumptions hidden here: the gauge $\nu(x)$ might not be equivalent to distance to a basepoint, one may need a family of uniformly continuous functions to account for a controlled function, and so on.

It is worth noting that the $\sup_x^{r,r'}$ constructions in [4] combined with the Riemann sum comment above can show how our $\sup_x]_{r_1}^{r_2}$ quantifiers can be built even in bounded versions of continuous logic.

2.14. **Pseudometrics.** The definitions and constructions in any version of continuous logic can be carried out extremely similarly when working with pseudometric spaces rather than metric spaces. That is, when we use sorts which are a set M or pointed set $(M, \star)$ equipped with a function $\mu : (x, y) \mapsto \mathbb{R}_{\geq 0}$ which is symmetric and satisfies the triangle inequality, but might have $\mu(x, y) = 0$ for some $x \neq y$. This is analogous to how, in discrete logic, things are very similar if one uses an equivalence relation rather than equality.

It became clear near the completion of this thesis that a logic could be developed like the present one, but where sorts are taken to be a set M equipped with a sequence of pseudometric-basepoint pairs $(\mu_i, \star_i)$. This would essentially be a logic for the much more general class of uniform spaces.

This would primarily require a slight extension of the definition of controlled functions. For example, the purpose of the N controller for a function $f : \prod_{n\in\alpha} M_n \to M$ would be to select finitely many pseudometrics from among the collection of all pseudometrics $\mu_{n,i}$ associated to one of the factors M_n in the domain of f. The other major change would come later in the ultraproduct construction, where additional boundedness requirements with respect to the pseudometrics would be necessary in order for the construction to remain in the correct class of spaces.

There are of course many details to be checked, but we suspect the development would closely mirror the development here. This project is not taken up in this thesis, but given the potential usefulness of the more general setting, it seems worthwhile to call attention to this observation.

3. Model Theory

We will now begin discussing topics having to do with the properties of structures and how structures relate to one another, as measured by the formulas of a given language.

3.1. **Theories and elementary maps.** We have seen that L-structures assign functions to L-formulas. In particular, L-structures assign a constant function to each L-sentence, and we can think of this constant function as just a real number. Interesting properties of M can be related to the values of sentences ϕ.

If we build a nicer structure N which we know agrees with M on ϕ, then we can learn about M by studying N instead. This line of thought is one of the fundamental insights of model theory.

Below, we define the theory of a structure. The theory of M summarizes the values of all sentences in M. Let us pause to consider the analogy with discrete logic. In the discrete setting, there are only two possibilities for the value of a sentence in: true or false. There, we can keep track of all values of all sentences in M by just noting which sentences are true. The false sentences are necessarily everything else, and this gives a complete account of the values of all sentences in M.

In the continuous setting, sentences take values in $\mathbb{R}$ rather than $\{\text{True}, \text{False}\}$. However, we will easily see that we still only need to keep track of the sentences corresponding to some fixed single value in $\mathbb{R}$. We will focus on those sentences with value 0. Of course, knowing that $\phi^M \neq 0$ does not tell us the value of ϕ^M. But if we have a complete list of all sentences ψ such that $\psi^M = 0$, then we can determine the value of ϕ^M by looking for which r satisfies $(\phi - r)^M = 0$, since $\phi - r$ will also be a sentence.

It is not particularly important that we use 0. We could choose any value for this purpose. But it is convenient to focus on the sentences which evaluate to 0. One advantage to focusing on 0 is that basic algebraic facts let us read these sentences in familiar, boolean-like ways. For a quick example, remember that $ab = 0$ iff $a = 0$ or $b = 0$. These are easy algebraic facts, but we will call attention to them in a later section. Also, focusing on sentences with value 0 means that theories will always contain things like $\sup_x]_r^{r'} d(x, x)$, which is the closest analog to a discrete sentence like $\forall x(x = x)$.

Definition 3.1.1. The **theory** of an L-structure M is denoted $\mathrm{Th}(M)$ and defined to be the collection of L-sentences ϕ such that $\phi^M = 0$. More generally, we may call any collection T of L-sentences a **theory** if there is at least one M such that $\phi^M = 0$ for all $\phi \in T$.

When we have identified some theory T, we are usually interested in talking about the structures whose theory contains T. So we have names and notation for this situation.

Definition 3.1.2. If ϕ is a sentence, we say M **satisfies** ϕ and write $M \models \phi$ whenever ϕ^M =0. If T is a collection of sentences, we say M is a **model** of T and write $M \models T$ whenever $\phi^M = 0$ for all $\phi \in T$.

Clearly, every structure is a model of its own theory. A more interesting case is when two different structures share the same theory.

Definition 3.1.3. Let M, N be L-structures. We say M and N are **elementarily equivalent** if $\mathrm{Th}(M) = \mathrm{Th}(N)$. Elementary equivalence of M and N is usually denoted by $M \equiv N$.

That is, M and N are elementarily equivalent exactly when they agree about which sentences are 0. We now check that elementary equivalence implies agreement of interpretation for all sentences, not just the 0 sentences.

Proposition 3.1.4. *Let M and N be two L-structures with* $\mathrm{Th}(M) = \mathrm{Th}(N)$. *Then for any L-sentence ϕ, we have $\phi^M = \phi^N$.*

Proof. Suppose $\phi^M = r$. Let ψ be the sentence $\phi - r$. Note that ψ is an L-sentence since we obtained it by composing ϕ with the connective $x \mapsto x - r$. Now, $\psi^M = 0$ by interpretation, and hence $\psi \in \mathrm{Th}(M)$. The hypothesis then gives $\psi^N = 0$, from which we get $\phi^N = r$ as well. □

Now we will define several important kinds of maps between L-structures.

Definition 3.1.5. Let M, N be L-structures. Suppose $\mathcal{E}$ is a collection of functions $\mathcal{E}_S : M_S \to N_S$ for each sort S in L. We will denote this as follows.

$$\mathcal{E} : M \to N$$

We say $\mathcal{E}$ is an **L-embedding** if it respects distances, functions and relations as follows:

- for each distance symbol d_S and $a, b \in M_S$,
$$d_S^N(\mathcal{E}_S(a), \mathcal{E}_S(b)) = d_S^M(a, b),$$
- for each function symbol f with $\mathrm{domain}(f) = (S_n : n \in \mathbb{N})$ and $\mathrm{range}(f) = S$, and each $a_n \in M_{S_n}$,
$$f^N(\mathcal{E}_{S_n}(a_n) : n \in \mathbb{N}) = \mathcal{E}_S(f^M(a_n : n \in \mathbb{N})),$$
- for each relation symbol R with $\mathrm{domain}(f) = (S_n : n \in \mathbb{N})$ and each $a_n \in M_{S_n}$,
$$R^N(\mathcal{E}_{S_n}(a_n) : n \in \mathbb{N}) = R^M(a_n : n \in \mathbb{N}),$$

We say $\mathcal{E}$ is **L-elementary** if for any L-formula $\phi(x_n : n \in \mathbb{N})$ with each x_n of sort S_n respectively, we have

$$\phi^N(\mathcal{E}(a_n) : n \in \mathbb{N}) = \phi^M(a_n : n \in \mathbb{N})$$

whenever $a_n \in M_{S_n}$.

We say $\mathcal{E}$ is an **L-isomorphism** if it is a surjective L-embedding. Isomorphism of M and N is usually denoted by $M \cong N$ when the language is understood.

The next proposition highlights basic relations between these kinds of maps. Informally, the difference between embeddings and elementary maps is that, when $\mathcal{E} : M \to N$ is only an embedding, there may be some elements of $N - \mathcal{E}(M)$ which can be indirectly referred to by quantifiers. That is, the embedded copy of an element $\mathcal{E}(a) \in N$ might relate to "new" things in N differently than a did to things just in M. The sup and inf quantifiers can detect this occuring. A simple example can be built by noticing that $\{0\}$ embeds in $\{0, 1\}$, but $\sup_x]_1^2 d(x, 0)$ is 0 in $\{0\}$ and is 1 in $\{0, 1\}$. Elementary maps avoid this by requiring the map to respect formulas, not just symbols. Isomorphisms avoid it by requiring surjectivity, so that there is nothing new in N.

Proposition 3.1.6. *Any isomorphism is elementary. Any elementary map is an embedding. Any embedding is an isometry.*

Proof. We will assume we have single-sorted structures, and functions and relations are unary operators. The general cases are similar, but this makes things more readable.

Suppose $\mathcal{E}$ is an embedding. Then by definition we have

$$d^N(\mathcal{E}(a), \mathcal{E}(b)) = d^M(a, b)$$

for all $a, b \in M$. This is the definition of isometry.

Suppose $\mathcal{E}$ is elementary. For the distance symbol d, notice $d(x, y)$ is a formula. Similarly, for any relation symbol R, $R(x)$ is a formula. So we immediately get both

$$d^N(\mathcal{E}(a), \mathcal{E}(b)) = d^M(a, b)$$

and

$$R^N(\mathcal{E}(a)) = R^M(a)$$

for all meaningful $a, b \in M$ by assumption. For any function symbol f, notice $d(f(x), y)$ is a formula. This means

$$d^N(f^N(\mathcal{E}(a)), \mathcal{E}(b)) = d^M(f^M(a), b)$$

for all $a, b \in M$ by assumption. But, for any $a \in M$, we have

$$d^M(f^M(a), f^M(a)) = 0$$

and hence using $b = f^M(a)$ above gives us

$$d^N(f^N(\mathcal{E}(a)), \mathcal{E}(f^M(a))) = 0.$$

This is equivalent to

$$f^N(\mathcal{E}(a)) = \mathcal{E}(f^M(a))$$

for any $a \in M$, so we are done.

Suppose $\mathcal{E}$ is an isomorphism. We can argue that $\mathcal{E}$ is elementary by induction on formulas. It is clear that if ϕ is quantifier free, we have

$$\phi^N(\mathcal{E}(a)) = \phi^M(a)$$

since ϕ^N is just a composition of functions with this property. So, we just need to handle the quantifier case. Suppose ϕ is of the form $\sup_x]_r^{r'} \psi(x, y)$ where ψ has the desired property. Then for any $b \in M$, we have that $\phi^N(\mathcal{E}(b))$ is

$$\int_r^{r'} \sup_{x \in B_\rho(N)} \psi^N(x, \mathcal{E}(b)) dr.$$

In particular, $\mathcal{E}$ is assumed to be a surjective isometry, so this integral is equivalent to

$$\int_r^{r'} \sup_{x \in B_\rho(M)} \psi^N(\mathcal{E}(x), \mathcal{E}(b))dr$$

where the sup now ranges over balls in M. By the inductive assumption on ψ, this is equivalent to

$$\int_r^{r'} \sup_{x \in B_\rho(M)} \psi^M(x, b)dr$$

which is $\phi^M(b)$. The more general forms of ϕ and the inf case are similar. □

Definition 3.1.7. Let M, N be L-structures. We say M is a **substructure** of N and write $M \subseteq N$ if for each sort S in L, we have $M_S \subseteq N_S$ and the inclusion maps $M_S \to N_S$ form an L-embedding. In this case, we also say N is an **extension** of M.

If moreover, this collection of inclusions $M_S \to N_S$ is elementary, then we say M is an **elementary substructure** of N and write $M \preceq N$. In this case, we also say N is an **elementary extension** of M.

Proposition 3.1.8. *If $M \preceq N$, then* $\mathrm{Th}(M) = \mathrm{Th}(N)$.

Proof. Sentences are a special case of the formulas preserved by the elementary inclusion map. □

We will close this section with the Tarski-Vaught test. This provides a sufficient condition for an elementary embedding to exist by exploiting the observation that obstructions to elementary embeddings are detected by quantifiers.

Proposition 3.1.9. ***(Tarski-Vaught Test)*** *Let M, N be L-structures with $M \subseteq N$. The following are equivalent.*

(1) $M \preceq N$

(2) *For all $r_1, r_2 \in \mathbb{R}$, all L-formulas $\phi(x, (y_n))$, and all sequences (a_n) from M, whenever we have*

$$\left(\inf_x]_{r_1}^{r_2} \phi(x, (a_n))\right)^N < r$$

for some $r \in \mathbb{R}$, then there is some $b \in B_{r_2}(M)$ such that

$$(\phi(b, (a_n)))^N < r.$$

(3) *Let D be a dense subset of $\mathbb{R}$, let F be a dense subset of L-formulas, and let A be a dense subset of M. For all $r_1, r_2 \in D$, all $\phi(x, (y_n))$ from F, and all sequences (a_n) over A, whenever we have*

$$\left(\inf_x]_{r_1}^{r_2} \phi(x, (a_n))\right)^N < r$$

for some $r \in \mathbb{R}$, then there is some $b \in B_{r_2}(M)$ such that

$$(\phi(b, (a_n)))^N < r.$$

Proof. The $(1 \to 2)$ direction is by definition of $\preceq$. The bound on the quantified formula carries over to M, and then we can select the witness $b \in M$ by considering the bound on the integral we get by interpretation.

For $(2 \to 1)$, just note that induction on complexity of L-formulas can be carried out, using $M \subseteq N$ to establish the base case, and (2) to deal with the inductive steps involving quantification. Notice that we have (2) for all pairs r_1, r_2. So when we want to compare $\inf_x]_{r_1}^{r_2}\phi$ in M and N, we can arbitrarily partition $[r_1, r_2]$ into subintervals $[r_1', r_2']$ where (2) applies to find witnesses in every ball of N.

The implication $(3 \to 2)$ follows by using density of F, continuity of formulas, and continuity of integrals with respect to their endpoints in order to verify the conditions in 2 to within ε for every $\varepsilon > 0$. □

Notice that in (3) it suffices to check enough formulas and parameters, in the sense that we only need to consider dense sets. In particular, since formulas are controlled, we only actually need to worry about finite tuples (a_n).

3.2. **Notation for common subformulas.** It becomes immediately apparent in applications that we want to assert equalities and inequalities using formulas or sentences in theories. For a simplified example with equality, consider the following.

Example 3.2.1. We might want to assert that two constant points a and b must be at distance 1 from each other in some theory, i.e. that $d(a, b) = 1$. We have the distance symbol in our language, but we do not have equality. Instead, we can use the sentence ϕ given by $|d(a, b) - 1|$. By definition, $\phi \in \mathrm{Th}(M)$ if and only if M evalutes $|d(a, b) - 1|$ to be 0. So we have that $\phi \in \mathrm{Th}(M)$ iff $d(a, b) = 1$.

In our setting, ϕ is not "true" or "false" in M; it is assigned a real number. Knowing $\phi \notin \mathrm{Th}(M)$ only tells us $\phi^M \neq 0$, but not what value ϕ^M is. Considering the definition of ϕ, a reasonable way to think of ϕ^M is as the (absolute value of the) error in the equality statement, i.e. how much $d(a, b)$ differs from 1.

Now a simple example with inequality.

Example 3.2.2. Suppose we want to assert $d(a, b) \leq 1$. Again, we do not have $\leq$ in our language, so this is not a sentence. But we can consider the sentence ϕ given by $\max(d(a, b) - 1, 0)$. We have $\max(d(a, b) - 1, 0) = 0$ iff $d(a, b) \leq 1$, so we have $\phi \in \mathrm{Th}(M)$ exactly when $d(a, b) \leq 1$ in M. More generally, we have that $\max(d(a, b) - 1, 0) = r$ iff $d(a, b) = 1 + r$. That is, ϕ^M can be considered the error in the $\leq$ statement, i.e. how much greater $d(a, b)$ is than 1.

Because these are very typical situations, we introduce the following notation.

Notation 3.2.3. Let L be any language, and let ϕ, ψ be L-formulas.

(1) $\phi \approx \psi$ is the formula $|\phi - \psi|$.
(2) $\phi \lesssim \psi$ is the formula $\max(\phi - \psi, 0)$.
(3) $\phi \gtrsim \psi$ is the formula $\max(\psi - \phi, 0)$.
(4) For each $\varepsilon > 0$, $\phi \approx_\varepsilon \psi$ is the formula $(\phi \approx \psi) \lesssim \varepsilon$, i.e.

$$\max(|\phi - \psi| - \varepsilon, 0).$$

The inclusion of absolute values for $\approx$ forces the formula to be nonnegative. This is convenient when working with quantifiers and theories.

Proposition 3.2.4. *The following hold*
(1) $(\phi \approx \psi)^M = 0$ *iff* $\phi^M = \psi^M$.
(2) $(\phi \lesssim \psi)^M = 0$ *iff* $\phi^M \leq \psi^M$.
(3) $(\phi \gtrsim \psi)^M = 0$ *iff* $\phi^M \geq \psi^M$.
(4) $(\phi \approx_\varepsilon \psi)^M = 0$ *iff* $|\phi^M - \psi^M| \leq \varepsilon$.

3.3. **Reading, writing, and arithmetic.** Provided $x, y \geq 0$, the equations $xy = 0$ and $x + y = 0$ correspond to boolean claims about the terms involved. We have $xy = 0$ iff at least one of $x = 0$ or $y = 0$ holds. When $x, y \geq 0$, we have $x + y = 0$ iff both $x = 0$ and $y = 0$ hold.

Similarly, still provided $x, y \geq 0$, the equation $\max(x, y) = 0$ holds iff both $x, y = 0$, and $\min(x, y) = 0$ holds iff at least one of x, y is 0.

As was suggested earlier, these facts are part of the motivation for defining $\mathrm{Th}(M)$ to be the set of sentences evaluating to 0.

Example 3.3.1. Consider the statement *"the distance from a to b is* 1*, or the distance from a to b is* 2*"*. In discrete logic, it is clear enough how one might express this; for example, $d(a, b) = 1 \vee d(a, c) = 2$. We have seen how to handle $=$, but now we need to handle $\vee$. Consider the formula ϕ given by the following.

$$(d(a, b) \approx 1) \cdot (d(a, c) \approx 2)$$

We have $\phi^M = 0$ iff at least one of the terms in the product is 0. Equivalently, we have $\phi^M = 0$ iff $d(a, b) = 1$ or $d(a, c) = 2$. We could similarly have used the formula ψ given by

$$\min\left(d(a, b) \approx 1, d(a, c) \approx 2\right).$$

Example 3.3.2. Consider the statement *"the distance from a to b is* 1*, and the distance from a to c is at most* 2*"*. Again, in discrete logic, we could express this as $d(a, b) = 1 \wedge d(a, c) \leq 1$. For continuous logic, consider the formula ϕ given by the following.

$$(d(a, b) \approx 1) + (d(a, c) \lesssim 1)$$

Then, since both terms in the sum are bounded below by 0, we have $\phi^M = 0$ iff both terms in the sum are 0. Equivalently, we have $\phi^M = 0$ iff $d(a, b) = 1$ and $d(a, c) \leq 1$. We could similarly have used the formula ψ given by

$$\max\left(d(a, b) \approx 1, d(a, c) \lesssim 1\right).$$

We will summarize these observations in the following proposition.

Proposition 3.3.3. *Let ϕ and ψ be L-sentences, and let M be an L-structure. These four statements are equivalent:*

- $M \models \phi$ *or* $M \models \psi$.
- $\phi^M = 0$ *or* $\psi^M = 0$.
- $(\phi \cdot \psi)^M = 0$.
- $(\min(\phi, \psi))^M = 0$.

When $\phi^M, \psi^M \geq 0$, these four statements are equivalent:

- $M \models \phi$ *and* $M \models \psi$.
- $\phi^M = 0$ *and* $\psi^M = 0$.
- $(\phi + \psi)^M = 0$.
- $(\max(\phi, \psi))^M = 0$.

3.4. **Chains.** When $M \subseteq N$, it is trivial that we can take the union of the sorts in M and N to form an L-structure (in this case, N) which contains both M and N. It is also easy to see that this generalizes to chains of substructures.

Definition 3.4.1. Let I be an ordered set. For all $i \in I$, let M_i be an L-structure, such that $M_i \subseteq M_j$ whenever $i < j$. The **union** $\bigcup_{i \in I} M_i$ is the L-structure formed by taking the union and metric completion along each sort, function, and relation.

If, moreover, the chain consists of elementary substructures, rather than just substructures, we have the following.

Proposition 3.4.2. *Let I be an ordered set, and let M_i be an L-structure for each $i \in I$. If $M_i \preceq M_j$ whenever $i < j$, then the union satisfies*

$$M_i \preceq \bigcup_{j \in I} M_j$$

for all $i \in I$.

Proof. By the Tarski-Vaught test, it suffices to only consider inf] quantified formulas. We can approximate such a formula with a Riemann sum to reduce to finding finitely many witnesses to inf statements. But any finitely tuple in $\bigcup M$ must be contained in a common M_i. $\square$

3.5. **Downward Löwenheim-Skolem.** Given a structure M and a subset A, we might want to construct an elementary substructure N which is smaller than M but still contains A. That is, we might want a smaller model which still respects some core behavior.

The next result explains when and how this can be done. The main idea is to start with A, and inductively add the elements needed in order to satisfy the Tarski-Vaught test. That is, we just need to add witnesses for things like $\inf_x]_{q_1}^{q_2} \phi(x, a)$. Since the language is determined by density(L) many formulas, we will not have to expand A by too much.

In the following, we write density$(A) \leq \kappa$ to mean that A has a subset B which is dense in A and whose cardinality is $\leq \kappa$.

Proposition 3.5.1. ***(Downward Löwenheim-Skolem)*** *Let κ be an infinite cardinal, let L be a language with* density$(L) \leq \kappa$. *Let M be an L-structure, and let $A \subseteq M$ with* density$(A) \leq \kappa$.

There exists an elementary substructure $N \preceq M$ such that $A \subseteq N$ and density$(N) \leq \kappa$.

Proof. Let F be a dense set of L-formulas such that each $\phi \in F$ has finitely many variables. Choose some dense subset B of A with cardinality $\leq \kappa$. Let $B_0 = B$, and let $\bar{B}_0$ be the closure (under functions from M) and completion of B_0. Note that since L has density $\leq \kappa$, we still have that the density of $\bar{B}_0$ is $\leq \kappa$, but now $\bar{B}_0$ can be taken to be an L-structure.

Now, for each $n \in \{1, 2, 3, \dots\}$, construct the "witness set" W_n as follows. For each $\phi(x_1, \dots, x_k, y)$ in F and each $b_1, \dots, b_k \in B_{n-1}$, whenever there is an $m \in \mathbb{N}$ and an element c satisfying

$$c \in B_{\frac{1}{n}(m+1)}(M) - B_{\frac{1}{n}(m)}(M)$$

and

$$M \models |\phi(b_1, \dots, b_k, c)| \lesssim \frac{1}{n},$$

we add one such c (for each m) to W_n. Note that this means we add at most countably many such c for each ϕ.

For $n \in \{1, 2, 3, \dots\}$, define B_n to be $B_{n-1} \cup W_n$, and define $\bar{B}_n$ to be the closure (under functions) and completion of B_n. Each B_n still has density $\leq \kappa$ since we assumed the density of L is $\leq \kappa$. So, the union of the $\bar{B}_n$ is now an L-structure of density $\leq \kappa$.

Moreover, for any $\phi \in F$ and $b_1, \dots, b_k \in \bigcup \bar{B}_n$, and for any $d \in M$, our construction ensures that we can always approximate

$$(\phi(b_1, \dots, b_k, d))^M$$

to within any $\varepsilon > 0$ by some

$$(\phi(b_1, \dots, b_k, c))^{\bigcup \bar{B}_n}$$

with $c \in \bigcup \bar{B}_n$ having $||c||$ within ε of $||d||$. This is sufficient to verify the inf] equalities needed in the Tarski-Vaught test and conclude that

$$A \subseteq \bigcup B_n \preceq M.$$

□

3.6. **Ultrafilters.** If we have a collection of L-structures $\{M_i : i \in I\}$, we would like to produce an L-structure M which ideally represents the tendencies of the collection. That is, we want a metric structure where each interpreted formula ϕ^M should in some way represent the collection $\{\phi^{M_i} : i \in I\}$.

Formalizing this begins with the idea of an ultrafilter, which gives us one way to define being a "sufficiently large" subset of I. A choice of ultrafilter is will give the collection of M_i some notion of direction or tendency.

Definition 3.6.1. An **ultrafilter** on a set I is a collection $\mathcal{U}$ of subsets of I such that:

(1) $\emptyset$ is not in $\mathcal{U}$.
(2) If $F_1 \in \mathcal{U}$ and F_2 is a subset of I with $F_1 \subseteq F_2$, then $F_2 \in \mathcal{U}$ as well.
(3) If $F_1, F_2 \in \mathcal{U}$, then the intersection $F_1 \cap F_2$ is in $\mathcal{U}$.
(4) If F is a subset of I, then either $F \in \mathcal{U}$ or $I \setminus F$ is in $\mathcal{U}$.

An ultrafilter is **nonprincipal** if it does not contain a singleton set.

Nonprincipal ultrafilters are far more interesting for our purposes. We will discuss limits with respect to ultrafilters, and if $\mathcal{U}$ contains a singleton $\{i\}$ (i.e. if $\mathcal{U}$ is principal), then these limits only depend on the single element i. This would be like defining a limit of real sequences x_n by choosing a specific coordinate of the sequence, rather than by considering the tails of the sequence.

As mentioned above, a good way to think of an ultrafilter is that it decides whether any subset of I is sufficiently large. In this sense, $\emptyset$ is not $\mathcal{U}$-large, anything containing a $\mathcal{U}$-large set is $\mathcal{U}$-large, and the intersection of $\mathcal{U}$-large sets is $\mathcal{U}$-large. The "ultra" part of the ultrafilter is the requirement that every set is either considered $\mathcal{U}$-large or else has a $\mathcal{U}$-large complement.

The next proposition is just a dual version of the intersection property above.

Proposition 3.6.2. *If none of $F_1, \dots, F_n$ are in $\mathcal{U}$, then $F_1 \cup \dots \cup F_n$ is not in $\mathcal{U}$.*

Proof. We have that the complements $I \setminus F_1, \dots, I \setminus F_n$ are each in $\mathcal{U}$. Thus their intersection $(I \setminus F_1) \cap \dots \cap (I \setminus F_n)$ is in $\mathcal{U}$. But this intersection is $I \setminus (F_1 \cup \dots \cup F_n)$. So, its complement $F_1 \cup \dots \cup F_n$ is not in $\mathcal{U}$. $\square$

It is straightforward from the last proposition that any nonprincipal ultrafilter on $\mathbb{N}$ must contain the cofinite sets. In particular it must contain $\{1, 2, 3, \dots\}$, $\{2, 3, 4, \dots\}$, $\{3, 4, 5, \dots\}$, and so on. This is useful to keep in mind when seeing how ultrafilters can be used to generalize the usual notions of a limit from analysis.

Definition 3.6.3. Let x_i be an I-indexed collection of points in a metric space (X, d), and let $\mathcal{U}$ be an ultrafilter on I. We say $p \in X$ is the **ultralimit** of x_i and write $\lim_{\mathcal{U}} x_i = p$ if for all $\varepsilon > 0$, there is $F \in \mathcal{U}$ such that for all $i \in F$, we have $d(x_i, p) \leq \varepsilon$.

The proposition below shows an equivalent way to define ultralimits.

Proposition 3.6.4. *For any $\varepsilon > 0$, the following are equivalent:*

(1) *There is $F \in \mathcal{U}$ such that for all $i \in F$, we have $d(x_i, p) \leq \varepsilon$.*
(2) *The set $\{i \in I : d(x_i, p) \leq \varepsilon\}$ is contained in $\mathcal{U}$.*

Proof. For (1) implies (2), note that the set in (2) is a superset of F. For the other direction, let F be the given set. $\square$

The next few propositions show the advantages of using ultralimits. Namely, ultralimits have very general convergence properties.

Proposition 3.6.5. *If (x_i) has an ultralimit $\lim_{\mathcal{U}} x_i$, the ultralimit is unique.*

Proof. For contradiction, suppose p_1 and p_2 are two distinct ultralimits of (x_i) with respect to $\mathcal{U}$. Take balls B_1 and B_2 centered at p_1, p_2 respectively, and having $B_1 \cap B_2 = \emptyset$. Then $\{i \in I : x_i \in B_1\} \in \mathcal{U}$ and $\{i \in I : x_i \in B_2\} \in \mathcal{U}$, but their intersection is $\emptyset$. This is impossible. $\square$

Proposition 3.6.6. *If (M, d) is compact, then $\lim_{\mathcal{U}} x_i$ exists in M for any sequence (x_i).*

Proof. For contradiction, suppose no $p \in M$ is an ultralimit of (x_i). This means that for each $p \in M$, there is some $\varepsilon_p > 0$ such that

$$\{i \in I : d(x, p) \leq \varepsilon_p\} \notin \mathcal{U}.$$

Cover M by the balls centered at each p of radius ε_p, then let $B_1, \ldots, B_n$ be the balls in a finite subcover. Now,

$$\{i \in I : x_i \in B_k\} \notin \mathcal{U}$$

for each k, so the complements satisfy

$$\{i \in I : x_i \notin B_k\} \in \mathcal{U}.$$

Since there are finitely many, their intersection satisfies

$$\{i \in I : x_i \notin B_1 \cup \cdots \cup B_n\} \in \mathcal{U}$$

and hence is nonempty. This implies the existence of points x_i not contained in $B_1 \cup \cdots \cup B_n$, contradicting that these sets form a cover. $\square$

We can easily generalize this to proper spaces if we also know $||x_i||$ is bounded on a $\mathcal{U}$-large set. In particular, this is useful later when we look at ultralimits of values of formulas.

Corollary 3.6.7. *Let $(M, d, \star)$ be a proper pointed metric space, and let (x_i) be a sequence in M such that for some $r \in \mathbb{R}_+$, the set*

$$F = \{i \in I : ||x_i|| \leq r\}$$

is in $\mathcal{U}$.

Then $\lim_{\mathcal{U}} x_i$ exists and is unique.

Proof. Since M is proper, the closed ball $\bar{B}_r(M)$ is compact. Consider the sequence (x'_i) defined as follows.

$$x'_i = \begin{cases} x_i & \text{when } i \in F \\ \star & \text{when } i \notin F \end{cases}$$

The last proposition applies to (x_i') as a sequence in $\bar{B}_r(M,p)$, so it has an ultralimit p. By definition, this means that for any $\varepsilon > 0$, there is $G \in \mathcal{U}$ such that $d(x_i', p) \leq \varepsilon$ for $i \in G$. But then $d(x_i, p) \leq \varepsilon$ for all i in the intersection $F \cap G \in \mathcal{U}$. So p is an ultralimit of (x_i). Uniqueness follows from the earlier proposition. □

This last corollary shows that, for example, the alternating sequence

$$(1, 0, 1, 0, \dots)$$

has an ultralimit once we choose an ultrafilter. For this sequence, the ultralimit can be decided by determining whether the ultrafilter contains the set of even indices, or its complementary set of odd indices. If we choose an arbitrary ultrafilter, then we will not know what limit to expect. Also, of course, more elaborate examples can be constructed. It is comforting to know the following, at least.

Proposition 3.6.8. *If $(x_n : n \in \mathbb{N})$ is a sequence in M which has a limit $\lim_{n\to\infty} x_n$ in the usual metric sense, then for any non-principal ultrafilter $\mathcal{U}$, we have*

$$\lim_{\mathcal{U}} x_n = \lim_{n\to\infty} x_n.$$

Proof. This is because any non-principal ultrafilter must contain the cofinite sets

$$\{n, n+1, n+2, \dots\}$$

for each n. The usual definition of convergence then ensures that for each $\varepsilon > 0$, there is a cofinite set of indices where $d(x_i, \lim_{n\to\infty} x_n) \leq \varepsilon$. □

3.7. **Ultraproducts.** In this section we now define ultraproducts of L-structures and prove a fundamental theorem about them. The definition makes heavy use of the ultralimits from the previous section. We will check after the definition that the construction we describe is well-defined, and the resulting object is actually an L-structure.

Since the construction is long, we will break it up into two parts. The first part constructs the sorts for the ultraproduct. The second part defines the functions and relations on the sorts.

Definition 3.7.1. Let M_i be a collection of L-structures indexed by I, and let $\mathcal{U}$ be an ultrafilter on I. Let S be a sort symbol from L, and let $(M_i)_S$ denote the corresponding sort in each structure M_i.

We define the **ultraproduct sort** $\prod_{\mathcal{U}}(M_i)_S$ below in several steps.

(1) Let

$$\prod_{i\in I}(M_i)_S$$

be the cartesian product of the sorts $(M_i)_S$. Recall that each $(M_i)_S$ has a basepoint $\star_S^{M_i}$ and metric $d_S^{M_i}$.

(2) Define M'_S to be the subset

$$\{(a_i) \in \prod_{i\in I}(M_i)_S : \lim_{\mathcal{U}} d_S^{M_i}(a_i, \star_S^{M_i}) < \infty\}.$$

This is the set of sequences whose corresponding sequence of coordinate-wise distances to the basepoint has finite ultralimit.

(3) Define a pseudometric d'_S on M'_S by

$$d'_S(a,b) = \lim_{\mathcal{U}} d_S^{M_i}(a_i, b_i)$$

for each $a = (a_i)$ and $b = (b_i)$ in M'_S. This is the ultralimit of the coordinate-wise distances.

(4) Define M_S to be the quotient of M'_S by d'_S. That is, M_S is the metric space $M'_S/\sim$ where $\sim$ is the equivalence relation given by

$$a \sim b \text{ iff } d'_S(a,b) = 0.$$

This identifies the sequences whose ultralimit of coordinate-wise distances is 0.

(5) Define d_S^M to be the resulting quotient metric $d'_S/\sim$.

(6) Define $\star_S^M$ to be the class $(\star_S^{M_i} : i \in I)/\sim$.

(7) The ultraproduct sort $\prod_{\mathcal{U}}(M_i)_S$ is the pointed metric space $(M_S, d_S^M, \star_S^M)$

Now we define the full ultraproduct.

Definition 3.7.2. Let M_i be a collection of L-structures indexed by I, and let $\mathcal{U}$ be an ultrafilter on I. The **ultraproduct** $\prod_{\mathcal{U}} M_i$ is the L-structure M defined as follows.

(1) For each sort symbol S in L, the sort M_S is the ultraproduct sort $\prod_{\mathcal{U}}(M_i)_S$ defined above.

(2) For each function symbol $f \in \mathcal{F}$ with domain$(f) = (S_n : n \in \alpha)$ and range$(f) = S$, we define

$$f^M : \prod_{n\in\alpha} M_{S_n} \to M_S$$

as follows. For each $(a_n : n \in \alpha) \in \prod M_{S_n}$, we define

$$f^M(a_n : n \in \alpha) = (f^{M_i}(a_{n,i} : n \in \alpha))_{i\in I}/\sim$$

where for each fixed n, the sequence $(a_{n,i})_{i\in I}$ is any representative for the class of a_n.

(3) For each relation symol $R \in \mathcal{R}$ with domain$(R) = (S_n : n \in \alpha)$, we define

$$R^M : \prod_{n\in\alpha} M_{S_n} \to \mathbb{R}$$

as follows. For each $(a_n : n \in \alpha) \in \prod M_{S_n}$, we define

$$R^M(a_n : n \in \alpha) = \lim_{\mathcal{U}} R^{M_i}(a_{n,i} : n \in \alpha)$$

where for each fixed n, the sequence $(a_{n,i})_{i\in I}$ is any representative for the class of a_n.

If all M_i are the same structure N, then we call $\prod_{\mathcal{U}} M_i$ the **ultrapower** of N, and denote it by $N^{\mathcal{U}}$.

Proposition 3.7.3. *The functions and relations above are well-defined and $\prod_{\mathcal{U}} M_i$ is an L-structure.*

Proof. We consider the case of functions. Relations follow from a similar argument. We need to check the existence and uniqueness of the value assigned to $f^M(a_n : n \in \alpha)$ in the construction above. Moreover, we need to verify that f^M is controlled by the tuple control$(f) = (\lambda, N, \delta)$ assigned by the language.

For each $n \in \alpha$, let $(a_{n,i})_{i\in I}$ be a representative for the class a_n, and let $r = (r_n) \in \mathbb{R}_+^\alpha$ be a real sequence with $||a_n|| < r_n$ for all $n \in \alpha$.

First, we will check that there is some $F \in \mathcal{U}$ where for each $i \in F$ the value of

$$f^{M_i}(a_{n,i} : n \in \alpha)$$

is bounded. This will verify that

$$(f^{M_i}(a_{n,i} : n \in \alpha))_{i\in I}$$

actually represents a class in M_S. So that we can apply controllers, choose any $\varepsilon > 0$. For all i, the function f^{M_i} is controlled by (λ, N, δ). For each $n \in N(r, \varepsilon)$, since a_n is represented by $(a_{n,i})_{i\in I}$, there is $F_n \in \mathcal{U}$ such that

$$||a_{n,i}|| < r_n \text{ for } i \in F_n.$$

Since $N(r, \varepsilon) \subseteq \alpha$ is finite, by intersecting these finitely many F_n, we get a single $F \in \mathcal{U}$ where for all $n \in N(r, \varepsilon)$ and $i \in F$ we have $||a_{n,i}|| < r_n$. Hence, for all $i \in F$, we have by definition of the controllers that

$$||f^{M_i}(a_{n,i} : n \in \alpha)|| \leq \lambda(r).$$

So $(f^{M_i}(a_{n,i} : n \in \alpha))_{i\in I}$ represents some class in the closed ball of radius $\lambda(r)$.

Next, we check that this class is independent of the choice of representatives for each a_n. For each n, let $(b_{n,i})_{i\in I}$ also be a representative for the class a_n. Let $\varepsilon > 0$. For each $n \in N(r, \varepsilon)$, since $(a_{n,i})_{i\in I}$ and $(b_{n,i})_{i\in I}$ are both in the class of a_n, there are $F_{a,n}$ and $F_{b,n}$ in $\mathcal{U}$ such that

$$||a_{n,i}|| < r_n \text{ for } i \in F_{a,n},$$
$$||b_{n,i}|| < r_n \text{ for } i \in F_{b,n},$$

and moreover there is $F_{a,b,n} \in \mathcal{U}$ such that

$$d^{M_i}(a_{n,i}, b_{n,i}) < \delta(r, \varepsilon) \text{ for } i \in F_{a,b,n}.$$

By intersecting the finitely many $F_{a,n}$, $F_{b,n}$, and $F_{a,b,n}$ with $n \in N(r, \varepsilon)$, we obtain a single F_ε such that for all $n \in N(r, \varepsilon)$ and all $i \in F_\varepsilon$, the above three inequalities hold. Hence, for all $i \in F_\varepsilon$, we have by definition of the controllers that

$$d(f^{M_i}(a_{n,i} : n \in \alpha), f^{M_i}(b_{n,i} : n \in \alpha)) \leq \varepsilon.$$

We can make this argument for any $\varepsilon > 0$, so we have that

$$(f^{M_i}(a_{n,i} : n \in \alpha))_{i \in I}$$

and

$$(f^{M_i}(b_{n,i} : n \in \alpha))_{i \in I}$$

represent the same class, since the ultralimit of their distances is 0.

This verifies the well-definedness of f^M. Now we will check that the controllers are preserved.

Let $r = (r_n) \in \mathbb{R}_+^\alpha$ and let $\varepsilon > 0$. Suppose that we have sequences

$$a = (a_n : n \in \alpha)$$

$$b = (b_n : n \in \alpha)$$

from the ultraproduct M, such that for all $n \in N(r, \varepsilon)$ we have

$$||a_n|| < r_n,$$

$$||b_n|| < r_n,$$

$$d(a_n, b_n) < \delta(r, \varepsilon).$$

For each $n \in \alpha$, let $(a_{n,i})_{i \in I}$ be a representative for a_n, and let $(b_{n,i})_{i \in I}$ be a representative for b_n.

For all $n \in N(r, \varepsilon)$, since $||a_n|| < r_n$, there is some $F_n \in \mathcal{U}$ such that

$$||a_{n,i}|| < r_n \text{ for } i \in F_n.$$

If we intersect these finitely many F_n, we get a single $F \in \mathcal{U}$ such that for all $n \in N(r, \varepsilon)$ and $i \in F$, we have $||a_{n,i}|| < r_n$. For $i \in F$, we thus have

$$||f^{M_i}(a_{n,i} : n \in \alpha)|| \leq \lambda(r).$$

So, taking the class shows $||f^M(a)|| \leq \lambda(r)$ as required.

For all $n \in N(r, \varepsilon)$, since $||a_n|| < r_n$, $||b_n|| < r_n$, and $d(a_n, b_n) < \delta(r, \varepsilon)$, we have $F_{a,n}$, $F_{b,n}$, and $F_{a,b,n}$ such that

$$||a_{n,i}|| < r_n \text{ for } i \in F_{a,n},$$

$$||b_{n,i}|| < r_n \text{ for } i \in F_{b,n},$$

$$d^{M_k}(a_{n,i}, b_{n,i}) < \delta(r, \varepsilon) \text{ for } i \in F_{a,b,n}.$$

Intersecting these finitely many $F_{a,n}$, $F_{b,n}$, and $F_{a,b,n}$ for n in $N(r, \varepsilon)$ gives a single $F_\varepsilon \in \mathcal{U}$ such that for all $n \in N(r, \varepsilon)$ and $i \in F_\varepsilon$, we have $||a_{n,i}|| < r_n$, $||b_{n,i}|| < r_n$, and $d^{M_i}(a_{n,i}, b_{n,i}) < \delta(r, \varepsilon)$. Hence, for $i \in F_\varepsilon$, we get

$$d(f^{M_i}(a_{n,i} : n \in \alpha), f^{M_i}(b_{n,i} : n \in \alpha)) \leq \varepsilon.$$

So, $d(f^M(a), f^M(b)) \leq \varepsilon$ as required. $\square$

We have arrived at the analog of Łoś's Theorem for this logic. It shows that the interpretation of any formula in the ultraproduct is determined by ultralimits over the factors. The ultraproduct is a fundamental construction in model theory, and this theorem along with the compactness theorem we obtain from it below are central.

Theorem 3.7.4. ***(Fundamental theorem of ultraproducts)*** *Let $M = \prod_{\mathcal{U}} M_i$ be an ultraproduct of L-structures, and let ϕ be an L-formula with free variables $(x_n : n \in \alpha)$. The functions ϕ^M and ϕ^{M_i} satisfy*

$$\phi^M(a) = \lim_{\mathcal{U}} \phi^{M_i}(a_i)$$

for all sequences $a = (a_n : n \in \alpha)$ over appropriate sorts of M and sequences $a_i = (a_{n,i} : n \in \alpha)$ over appropriate sorts of M_i, where for each $n \in \alpha$, $(a_{n,i})_{i \in I}$ represents the class a_n in M.

In particular, if ϕ has no free variables, then the constant real value of ϕ^M is equal to the ultralimit of those for the ϕ^{M_i}.

Proof. The proof is by induction on ϕ.

First, we check that for any L-term t, we have

$$t^M(a) = (t^{M_i}(a_i) : i \in I)/\sim .$$

This is true by definition for $\star^M$, and true for variables since at any element $b \in M$, it is true by construction of the ultraproduct. Suppose it is true for all t_n with $n \in \alpha$, and let f be a function symbol. By interpretation,

$$f(t_n : n \in \alpha)^M(a)$$

is

$$f^M(t_n^M : n \in \alpha)(a).$$

By the inductive hypothesis, this is the same as

$$f^M((t_n^{M_i}(a_i) : i \in I)/\sim) : n \in \alpha).$$

By definition of the function f^M in the ultraproduct, this is

$$(f^{M_i}(t_n^{M_i}(a_i) : n \in \alpha) : i \in I)/\sim$$

as required.

Suppose ϕ is an atomic formula

$$R(t_n : n \in \alpha).$$

By interpretation ϕ^M is

$$R^M(t_n^M : n \in \alpha)$$

and so $\phi^M(a)$ is

$$R^M((t_n^{M_i}(a_i) : i \in I)/\sim) : n \in \alpha)$$

by the above paragraph on terms. By definition of the relation R^M in the ultraproduct, this means $\phi^M(a)$ is

$$\lim_{\mathcal{U}} R^{M_i}(t_n^{M_i}(a_i) : n \in \alpha).$$

Finally, by definition of interpretation in each M_i, this is the same as

$$\lim_{\mathcal{U}} R(t_n : n \in \alpha)^{M_i}(a_i)$$

as required.

Suppose ϕ is of the form

$$u(\psi_n : n \in \alpha)$$

where u is controlled on $\mathbb{R}^{\mathbb{N}}$ and the theorem holds for each ψ_n. By interpretation, $\phi^M(a)$ is

$$u(\psi_n^M(a) : n \in \alpha).$$

We want to show that this is the same $\lim_{\mathcal{U}} u(\psi_n^{M_i}(a_i) : n \in \alpha)$ to verify this case of the theorem. Let $r = (r_n) \in \mathbb{R}^\alpha$ be a sequence such that

$$|\psi_n^M(a)| < r_n$$

for all $n \in \alpha$, and let $\varepsilon > 0$. Choose controllers (λ, N, δ) for u. By the inductive hypothesis on each ψ_n, we have

$$\psi_n^M(a) = \lim_{\mathcal{U}} \psi_n^{M_i}(a_i)$$

for all $n \in \alpha$. So for each $n \in N(r, \varepsilon)$, there is $F_n \in \mathcal{U}$ such that

$$\left|\psi_n^{M_i}(a_i) - \psi_n^M(a)\right| < \delta(r, \varepsilon) \text{ for } i \in F_n$$

and

$$|\psi_n^{M_i}(a_i)| < r_n \text{ for } i \in F_n.$$

Intersecting these finitely many F_n with $n \in N(r, \varepsilon)$ gets us a single $F_\varepsilon \in \mathcal{U}$ such that for all $i \in F_\varepsilon$, we have by definition of the controllers that

$$\left|u\left(\psi_n^{M_i}(a_i) : n \in \alpha\right) - u\left(\psi_n^M(a) : n \in \alpha\right)\right| \le \varepsilon.$$

This shows

$$u\left(\psi_n^M(a) : n \in \alpha\right) = \lim_{\mathcal{U}} u\left(\psi_n^{M_i}(a_i) : n \in \alpha\right)$$

as required.

Suppose ϕ is of the form $\inf_x]_{r_1}^{r_2}\psi(x, y)$, where the theorem holds for ψ. We want to show that

$$\lim_{\mathcal{U}} \int_{r_1}^{r_2} \inf_{x \in \bar{B}_\rho(M_i)} \psi^{M_i}(x, a_i) d\rho = \int_{r_1}^{r_2} \inf_{x \in \bar{B}_\rho(M)} \psi^M(x, a) d\rho.$$

Let A be the value of the left-hand side, and let B be the value of the right-hand side. We will prove the theorem by showing that for any $L \in \mathbb{R}$,

$$A < L \implies B \le L$$

and

$$B < L \implies A \le L.$$

First, suppose $A < L$. Then there is some $F_1 \in \mathcal{U}$ such that for all $i \in F_1$, we have

$$\int_{r_1}^{r_2} \inf_{x \in \bar{B}_\rho(M_i)} \psi^{M_i}(x, a_i) dr < L.$$

Let $\varepsilon > 0$. We will select witnesses close to the infimums in each $\bar{B}_\rho(M_i)$ as we vary ρ, i, and use these witnesses to establish bounds on the infumums in the ultraproduct. For each $\rho \in [r_1, r_2]$ and $i \in I$, choose some $w_{\rho,i} \in \bar{B}_\rho(M_i)$ satisfying

$$\psi^{M_i}(w_{\rho,i}, a_i) < \inf_{x \in \bar{B}_\rho(M_i)} \psi^{M_i}(x, a_i) + \varepsilon.$$

Taking ultralimits on both sides of this inequality and writing w_ρ for each $(w_{\rho,i})/\sim$, we get

$$\psi^M(w_\rho, a) \leq \lim_{\mathcal{U}} \inf_{x \in \bar{B}_\rho(M_i)} \psi^{M_i}(x, a_i) + \varepsilon.$$

Since $w_{\rho,i} \in \bar{B}_\rho(M_i)$ holds for all ρ and i, we have each $w_\rho \in \bar{B}_\rho(M)$. So, we get the inequality

$$\inf_{x \in \bar{B}_\rho(M)} \psi^M(x, a) \leq \psi^M(w_\rho, a).$$

By transitivity, we have

$$\inf_{x \in \bar{B}_\rho(M)} \psi^M(x, a) \leq \lim_{\mathcal{U}} \inf_{x \in \bar{B}_\rho(M_i)} \psi^{M_i}(x, a_i) + \varepsilon.$$

Now, note that there is an $F_2 \in \mathcal{U}$ where for all $i \in F_2$ we have

$$\lim_{\mathcal{U}} \inf_{x \in \bar{B}_\rho(M_i)} \psi^{M_i}(x, a_i) \leq \inf_{x \in \bar{B}_\rho(M_i)} \psi^{M_i}(x, a_i) + \varepsilon.$$

Combining the last two inequalities yields for all $i \in F_2$ that

$$\inf_{x \in \bar{B}_\rho(M)} \psi^M(x, a) \leq \inf_{x \in \bar{B}_\rho(M_i)} \psi^{M_i}(x, a_i) + 2\varepsilon.$$

By integrating this inequality over $\rho \in [r_1, r_2]$, we get for all $i \in F_2$ that

$$\int_{r_1}^{r_2} \inf_{x \in \bar{B}_\rho(M)} \psi^M(x, a) dr \leq \int_{r_1}^{r_2} \Big(\inf_{x \in \bar{B}_\rho(M_i)} \psi^{M_i}(x, a_i) + 2\varepsilon \Big) dr.$$

Finally, letting $F_\varepsilon = F_1 \cap F_2$, we have found a single $F_\varepsilon \in \mathcal{U}$ such that for all $i \in F_\varepsilon$ we have the following.

$$\begin{aligned} B &= \int_{r_1}^{r_2} \inf_{x \in \bar{B}_\rho(M)} \psi^M(x, a) dr \\ &\leq \int_{r_1}^{r_2} \inf_{x \in \bar{B}_\rho(M_i)} \psi^{M_i}(x, a_i) dr + 2\varepsilon \int_{r_1}^{r_2} 1 dr \\ &< L + 2\varepsilon \cdot |r_2 - r_1| \end{aligned}$$

This holds for all $\varepsilon > 0$, so we have shown $B \leq L$. This establishes the first part.

Now, suppose $B < L$. Let $\varepsilon > 0$. This argument is easier if we use the open ball version of the integrals. That is, we have assumed

$$B = \int_{r_1}^{r_2} \inf_{x \in B_\rho(M)} \psi^M(x, a) d\rho < L.$$

Let $R = (R_n) \in \mathbb{R}_+^{\mathbb{N}}$ be such that $||a|| < R_n$ for all n. Choose an (L, R, ε)-good partition $\rho_0 < \cdots < \rho_{K-1}$ for (the formula for) B on $[r_1, r_2]$, as in 2.7.6. This partition is such that B is strictly within ε of the Riemann sum

$$\Delta \cdot \sum_{k<K} \inf_{x \in B_{\rho_k}(M)} \psi^M(x, a).$$

Moreover, as long as we have $||a_{n,i}|| < R_n$ for some fixed finite number of coordinates, we will also have

$$\left| \int_{r_1}^{r_2} \inf_{x \in B_\rho(M_i)} \psi^{M_i}(x, a_i) d\rho - \Delta \cdot \sum_{k<K} \inf_{x \in B_{\rho_k}(M_i)} \psi^M(x, a_i) \right| < \varepsilon.$$

Let $G \in \mathcal{U}$ be such that $||a_{n,i}|| < R_n$ for $i \in G$ and for the finitely many required n. For each $k < K$, choose some witness $w_k \in B_{\rho_k}(M)$ satisfying

$$\psi^M(w_k, a) < \inf_{x \in B_{\rho_k}(M)} \psi^M(x, a) + \varepsilon \Delta^{-1} K^{-1}.$$

Fix a choice of representative $(w_{k,i})_{i \in I}$ for each w_k. By the inductive hypothesis on ψ, for each $k < K$, there is some $F_{1,k} \in \mathcal{U}$ where for all $i \in F_{1,k}$ we have

$$\psi^{M_i}(w_{k,i}, a_i) < \inf_{x \in B_{\rho_k}(M)} \psi^M(x, a) + \varepsilon \Delta^{-1} K^{-1}.$$

For each $k < K$ we can also find an $F_{2,k} \in \mathcal{U}$ where for all $i \in F_{2,k}$ we have $||w_{k,i}|| < \rho_k$ and hence $w_{k,i} \in B_{\rho_k}(M_i)$. Let F be the intersection of G with the finitely many $F_{1,k}$ and $F_{2,k}$ having $k < K$. Then $F \in \mathcal{U}$ and for all $i \in F$ and $k < K$, we have

$$\inf_{x \in B_{\rho_k}(M_i)} \psi^{M_i}(x, a_i) \leq \psi^{M_i}(w_{k,i}, a_i)$$

and so by transitivity

$$\inf_{x \in B_{\rho_k}(M_i)} \psi^{M_i}(x, a_i) < \inf_{x \in B_{\rho_k}(M)} \psi^M(x, a) + \varepsilon \Delta^{-1} K^{-1}.$$

Take the Riemann sum over both sides to get the following for all $i \in F$.

$$\begin{aligned} \Delta \cdot \sum_{k<K} \inf_{x \in B_{\rho_k}(M_i)} \psi^{M_i}(x, a_i) &< \Delta \cdot \sum_{k<K} \left(\inf_{x \in B_{\rho_k}(M)} \psi^M(x, a) + \varepsilon \Delta^{-1} K^{-1} \right) \\ &= \Delta \cdot \sum_{k<K} \left(\inf_{x \in B_{\rho_k}(M)} \psi^M(x, a) \right) + \varepsilon \\ &< B + 2\varepsilon \\ &< L + 2\varepsilon \end{aligned}$$

Since the partition was (L, R, ε)-good, we also have for all i that

$$\Delta \cdot \sum_{k<K} \inf_{x \in B_{\rho_k}(M_i)} \psi^{M_i}(x, a_i)$$

is strictly within ε of

$$\int_{r_1}^{r_2} \inf_{x \in B_\rho(M_i)} \psi^{M_i}(x, a_i).$$

So by combining this with the above set of inequalities, we have the following for all $i \in F$.

$$\begin{aligned} \int_{r_1}^{r_2} \inf_{x \in B_\rho(M_i)} \psi^{M_i}(x, a_i) &< \Delta \cdot \sum_{k<K} \inf_{x \in B_{\rho_k}(M_i)} \psi^{M_i}(x, a_i) \\ &< L + 3\varepsilon \end{aligned}$$

Taking an ultralimit then shows the following.

$$\begin{aligned} A &= \lim_{\mathcal{U}} \int_{r_1}^{r_2} \inf_{x \in B_\rho(M_i)} \psi^{M_i}(x, a_i) \\ &\leq L + 3\varepsilon \end{aligned}$$

This holds for all $\varepsilon > 0$, so we have $A \leq L$.

This completes the proof. The case where ϕ is a $\sup_x]$ formula is similar, or can be obtained via the fact that $\sup(x) = -\inf(-x)$. □

Corollary 3.7.5. ***(The Compactness Theorem)*** *Let T be a collection of L-sentences. The following are equivalent.*

(1) *For each finite $F \subseteq T$ and each $\varepsilon > 0$, there is an L-structure M such that*

$$|\phi^M| \leq \varepsilon \textit{ for all } \phi \in F.$$

(2) *For each finite $F \subseteq T$, there is an L-structure M such that*

$$M \models \phi \textit{ for all } \phi \in F.$$

(3) *There is a model $M \models T$.*

Proof. $(3 \to 2)$ and $(2 \to 1)$ are both clear, so we only need to check $(1 \to 3)$.

Consider the set of pairs

$$I = \{(F, \varepsilon) : F \text{ is a finite subset of } T, \text{ and } \varepsilon > 0\}.$$

For each $(F, \varepsilon) \in I$, let $M_{(F,\varepsilon)}$ be an L-structure with $|\phi^{M_{(F,\varepsilon)}}| \leq \varepsilon$ for each $\phi \in F$. We construct an ultrafilter $\mathcal{U}$ on I as follows. For each $(F, \varepsilon) \in I$, let

$$\uparrow (F, \varepsilon) \downarrow$$

denote the subset of I given by

$$\{(F', \varepsilon') \in I : F \subseteq F' \text{ and } \varepsilon' \leq \varepsilon\}.$$

The collection J of all such subsets $\uparrow (F, \varepsilon) \downarrow$ has the finite intersection property, i.e. any finitely many $\uparrow (F_n, \varepsilon_n) \downarrow$ have at least the element $(\bigcup_n F_n, \min_n \varepsilon_n)$ in common. So we can construct an ultrafilter $\mathcal{U}$ on I containing the collection J.

Now, consider the ultraproduct $M = \prod_{\mathcal{U}} M_{(F,\varepsilon)}$, and suppose $\phi \in T$. Then $\{\phi\}$ is a finite subset of T, and so for each $\varepsilon > 0$, we have $(\{\phi\}, \varepsilon) \in I$ and $\uparrow (\{\phi\}, \varepsilon) \downarrow$ in $\mathcal{U}$. But we know that

$$|\phi^{M_{(F',\varepsilon')}}| \leq \varepsilon$$

for all $(F', \varepsilon') \in \uparrow (\{\phi\}, \varepsilon) \downarrow$. So by definition of ultralimits, we get

$$\lim_{\mathcal{U}} \phi^{M_{(F,\varepsilon)}} = 0.$$

The fundamental theorem of ultraproducts yields $\phi^M = 0$. □

In discrete logic and in the bounded case for continuous logic, the compactness theorem can be applied to sets of formulas as well by adding constants to the language and working with the corresponding sentences in the larger language. In our setting, the addition of a constant is achieved by adding a 0-ary function symbol c. This requires specifying controllers, and notably the controller λ effectively acts as a bound on $||c||$. Consequently, the "formula version" of the compactness theorem in this setting only holds ball-wise. This becomes relevant later when we discuss the topology on type spaces.

As expected, structures embed elementarily into their ultrapowers.

Corollary 3.7.6. *Let $\mathcal{U}$ be an ultrafilter on I. For any ultrapower $M^{\mathcal{U}}$, the map $M \to M^{\mathcal{U}}$ defined by $a \mapsto (a : i \in I)/\sim$ is an elementary embedding.*

Proof. For all $a \in M$ we have the following.

$$\begin{aligned} \phi^M(a) &= \lim_{\mathcal{U}} \phi^M(a) \\ &= \phi^{M^{\mathcal{U}}}((a : i \in I)/\sim) \end{aligned}$$

The first equality is by definition of ultralimits. The second is by the fundamental theorem of ultraproducts. □

The following proposition indicates that proper structures are the analog in this setting of finite structures in discrete model theory, or compact structures in bounded continuous logic. Recall that bounded, proper structures are compact.

Proposition 3.7.7. *Let M be an L-structure in which each sort is proper. Then $M^{\mathcal{U}}$ and M are L-isomorphic for any ultrafilter $\mathcal{U}$.*

Proof. We will consider the single-sorted case, the more general case being clear from the argument.

Let $a \in M^{\mathcal{U}}$, and let $(a_i : i \in I)$ represent a. Since M is proper, $\lim_{\mathcal{U}} a_i$ exists when considered as a sequence in M. We will denote this ultralimit in M by $\bar{a}$. This ultralimit is independent of the choice of representative sequence for a, since any pair of such sequences have ultralimit distances going to 0. Thus, there is a function mapping every point $a \in M^{\mathcal{U}}$ to a corresponding point $\bar{a}$ in M.

This function is injective since if $\bar{a} = \bar{b}$, then $d(a, b)$ must be 0 in $M^{\mathcal{U}}$. Now, since $(\bar{a}, \bar{a}, \dots)$ represents a, the fundamental theorem of ultraproducts tells us that $\phi^{M^{\mathcal{U}}}(a) = \phi^M(\bar{a})$ for any $a \in M^{\mathcal{U}}$ and formula ϕ. The function is also clearly onto since we can obtain any element $b \in M$ as the image of $(b, b, \dots)/\sim$. So, we have an isomorphism. $\square$

3.8. **Elementary Classes.** Model theory is often concerned with the collection of models of some given theory T and the relations and maps between these models.

Definition 3.8.1. A class $\mathcal{C}$ of L-structures is called L-**axiomatizable** and also an L-**elementary class** if there is a theory T such that $M \models T$ if and only if $M \in \mathcal{C}$. In this case, the sentences in T are called **axioms** for $\mathcal{C}$.

One way to prove that we have an elementary class is to explicitly write down axioms and check the condition above. There is also a characterization of elementary classes in terms of ultraproducts.

Proposition 3.8.2. *Let $\mathcal{C}$ be a class of L-structures. The following are equivalent.*

(1) *$\mathcal{C}$ is an elementary class.*
(2) *$\mathcal{C}$ is closed under isomorphisms, ultraproducts, and elementary substructures.*

Proof. $(1 \to 2)$ follows from the definition of isomorphisms and elementary substructures, and the fundamental theorem of ultraproducts, which in particular tells us that values of sentences are preserved by ultraproducts.

We will now comment on $(2 \to 1)$. Let T be the set of sentences ϕ such that $\phi^M = 0$ for all $M \in \mathcal{C}$. This is our candidate set of axioms for the class. Clearly every $M \in \mathcal{C}$ satisfies T, so we just need to check the converse. Suppose $M \models T$.

Suppose (in order to obtain a contradiction) that some finite subset F of $\mathrm{Th}(M)$ were not satisfiable by a structure in $\mathcal{C}$. That is, there are sentences $\phi_1, \dots, \phi_n$ such that $\phi_i^M = 0$ for $i = 1, \dots, n$, but no $N \in \mathcal{C}$ has $\phi_i^N = 0$ for all i. There must be some $\varepsilon > 0$ such that all $N \in \mathcal{C}$ have $|\phi_i^N| \geq \varepsilon$ for at least one i, since otherwise we could take ultraproducts inside $\mathcal{C}$ to get an N with $\phi_i^N = 0$ for all i. By definition, T must contain the sentence $\max(\phi_1, \dots, \phi_n) \gtrsim \varepsilon$. But this contradicts $M \models T$. So, we conclude that $\mathrm{Th}(M)$ is finitely satisfiable within $\mathcal{C}$.

Now, again using closure of $\mathcal{C}$ under ultraproducts, we construct $N \models \mathrm{Th}(M)$ with $N \in \mathcal{C}$ by taking an ultraproduct along structures in $\mathcal{C}$ realizing an increasing chain of finite subsets of $\mathrm{Th}(M)$. So we have an $N \in \mathcal{C}$ with $N \equiv M$. By taking a large ultrapower of N, we obtain a model $N^{\mathcal{U}} \models \mathrm{Th}(M)$ which M elementarily embeds into. The closure properties of $\mathcal{C}$ then imply that $M \in \mathcal{C}$. $\square$

3.9. **Definable sets.** In any given L-structure, the zero sets of L-formulas determine subsets of the structure. Some of these zero sets are special in the sense that they behave like sorts. For example, we can sometimes quantify over a subset. These special zero sets are called definable sets, and are an important aspect of the structure.

Definability in continuous logic is less straightforward than in discrete logic. In general, it is not sufficient to consider level sets of formulas. For example, a set of the form

$$\{x \in M : \phi^M(x) = 0\}$$

will not generally be a definable set. We have seen this problem already when we discussed quantification over balls. The following example illustrates this problem a bit more abstractly. We can think of this as a quantification issue or as an issue with preserving zero sets under ultraproducts.

Example 3.9.1. Let L be a single-sorted language with relation symbols P and Q each controlled by (λ, N, δ) given by

$$(r \mapsto 1, (r, \varepsilon) \mapsto \emptyset, (r, \varepsilon) \mapsto \varepsilon).$$

Let M be the L-structure with sort $\mathbb{N}$ with the discrete metric. Define the relation P^M by

$$P^M(n) = \begin{cases} 0 & \text{when } n = 0 \\ \frac{1}{n} & \text{otherwise} \end{cases}.$$

Define Q^M by

$$Q^M(n) = \begin{cases} 0 & \text{when } n = 0 \\ 1 & \text{otherwise} \end{cases}.$$

Both are controlled as required since the metric is discrete.

We will see that there is no formula in L which interprets in models of $\mathrm{Th}(M)$ as the supremum of Q over elements of the zero set of P. Notice that

$$\{x \in M : P^M(x) = 0\}$$

is just $\{0\}$, and that $Q^M(0) = 0$. However, in the ultrapower $M^{\mathcal{U}}$, the set

$$\{x \in M^{\mathcal{U}} : P^{M^{\mathcal{U}}}(x) = 0\}$$

contains not just

$$\bar{0} = (0, 0, 0, \dots)/\sim$$

which has $Q^{M^{\mathcal{U}}}(\bar{0}) = 0$, but also more points. For example, it also contains the point $a = (n)_{n \in \mathbb{N}}/\sim$ which has $Q^{M^{\mathcal{U}}}(a) = 1$. This shows that M and $M^{\mathcal{U}}$ realize different supremums for values of Q over the zero set of P. So, the existence of a formula for this supremum would contradict Łoś's Theorem.

We know that M and $M^{\mathcal{U}}$ are elementarily equivalent, so this shows there cannot be a fixed sentence in L which interprets as the supremum of Q over the zero set of P. This phenomenon happens because $P^M(n)$ can tend to 0 without n approaching $\{x \in M : P^M(x) = 0\}$ in the metric.

Another notable departure from discrete logic is that in our setting we will have a notion of definable sets including not just finite tuples, but also sequences.

We will base our definition of definable sets on being able to approximate distance functions. This is useful for some of the upcoming proofs. However, we will give equivalent characterizations later which are more practical for checking that a given set is definable.

Definition 3.9.2. Let M be an L-structure, and let C be a set of elements from sorts of M.

If A is a subset of a finite product $\prod_{n=1}^N M_n$ of sorts, we say A is L-**definable** in M over C if for all positive real tuples

$$\begin{aligned} R &= (R_n : n \in \{1, \dots, N\}), \\ r &= (r_n : n \in \{1, \dots, N\}), \\ r' &= (r'_n : n \in \{1, \dots, N\}) \end{aligned}$$

with each $r'_n > r_n$, there is a sequence c over C and an L-formula $\phi_{R,r,r'}(x, z)$ such that $\phi^M_{R,r,r'}(x, c)$ is equivalent to the function

$$\left(\prod_{n=1}^{N} \frac{1}{|r'_n - r_n|}\right) \int_{r_N}^{r'_N} \cdots \int_{r_1}^{r'_1} d_N(x, \bar{B}_{(\rho_1,\dots,\rho_N)}(A)) d\rho_1 \cdots d\rho_N$$

when restricted to $x \in \bar{B}_{(R_1,\dots,R_N)}(M)$, where d_N is the max of the distances over the finitely many coordinates.

If A is a subset of an infinite product $\prod_{n\in\alpha} M_n$ of sorts, we say A is L-**definable** in M **over** C if for all finite subsets $F \subseteq \alpha$, the set $\pi_F(A)$ is L-definable in M over C.

We immediately have that any product of sorts is definable according to this definition. The formulas

$$\inf_y]_r^{r'} \max\left(d(x_1, y_1), \dots, d(x_N, y_N)\right)$$

for N-tuples y interpret as the required functions by definition of the inf] quantifier.

Readers familiar with bounded continuous logic will wonder how this definition relates to having formulas for $d(x, A)$. The following proposition clarifies this for the case where A is a subset of a finite product. We will continue writing d_N to mean the maximum of the distances in the coordinates.

Proposition 3.9.3. *Let A be a subset of a finite product $M = \prod_{n=1}^N M_n$ of sorts. Then we have the following.*

(1) *If A is bounded, i.e. $A \subseteq \bar{B}_{(R_1,\dots,R_N)}(M)$ for some $R_1, \dots, R_N$, then*
 (a) *for any r, r' with $r'_n > r_n \geq R_n$ for each $n \leq N$, the function*

$$\left(\prod_{n=1}^{N} \frac{1}{|r'_n - r_n|}\right) \int_{r_N}^{r'_N} \cdots \int_{r_1}^{r'_1} d_N(x, \bar{B}_{(\rho_1,\dots,\rho_N)}(A)) d\rho_1 \cdots d\rho_N$$

is the distance function $d_N(x, A)$

(b) *A is definable if and only if there is an L-formula which interprets as the function* $d_N(x, A)$.

(2) *If A is unbounded and has a point $a \in \bar{B}_{(R_1,\ldots,R_N)}(A)$, then*

(a) *for each real tuple $(r_1, \ldots, r_N)$, the function*

$$\int_{2r_N+R_N}^{2r_N+R_N+1} \cdots \int_{2r_1+R_1}^{2r_1+R_1+1} d(x, \bar{B}_{(\rho_1,\ldots,\rho_N)}(A))d\rho$$

is equivalent to $d_N(x, A)$ when restricted to $x \in \bar{B}_{(r_1,\ldots,r_N)}(M)$,

(b) *A is definable if and only if for each real tuple $(r_1, \ldots, r_N)$, there is an L-formula which interprets as a function equivalent to $d_N(x, A)$ when restricted to $x \in \bar{B}_{(r_1,\ldots,r_N)}(M)$.*

Proof. $(1a)$ Whenever $\rho_n \geq R_n$ for $n \leq N$, we have

$$d_N(x, \bar{B}_{(\rho_1,\ldots,\rho_N)}(A)) = d_N(x, A).$$

So the averages are over constant values equal to $d_N(x, A)$.

$(1b, \rightarrow)$ If A is definable, each

$$\left(\prod_{n=1}^{N} \frac{1}{|r'_n - r_n|}\right) \int_{r_N}^{r'_N} \cdots \int_{r_1}^{r'_1} d_N(x, \bar{B}_{(\rho_1,\ldots,\rho_N)}(A))d\rho_1 \cdots d\rho_N$$

is a formula by assumption. Applying $(1a)$ tells us these give us a formula for $d_N(x, A)$.

$(2a)$ We will discuss the single-variable case. Let $a \in A$, and let R be $d(a, \star)$. Then for any x and $r \in \mathbb{R}$, if we have $||x|| \leq r$, then we have

$$d(x, A) \leq r + R$$

by the triangle inequality, and hence

$$d(x, A) = d(x, \bar{B}_{r+(r+R)+\varepsilon}(A))$$

for any $\varepsilon > 0$. So for $||x|| \leq r$, we must have

$$d(x, A) = \int_{2r+R}^{2r+R+1} d(x, \bar{B}_\rho(A))d\rho$$

since this integral is just the average over the constant $d(x, A)$.

$(2b, \rightarrow)$ If A is definable, we have formulas equivalent to the integrals in $(2a)$ for $x \in \bar{B}_{(r_1,\ldots,r_N)}(M)$ by assumption. So we can apply $(2a)$ to get the the claimed L-formulas.

$(1b, 2b, \leftarrow)$ Suppose we want to obtain

$$\int_r^{r'} d(x, \bar{B}_\rho(A))d\rho$$

as a formula. We will assume $r < r'$. The hypotheses let us assume $d(y, A)$ is a formula, at least for $||y|| \leq R_y$ with R_y chosen larger than r', say $R_y > r' + 1$.

For any $\rho \in [r, r']$ and any $\varepsilon > 0$ less than 1, let $\phi_{\rho,\rho+\varepsilon}$ be a formula

$$\inf_y]_\rho^{\rho+\varepsilon} d(x,y) + (K/\varepsilon) \cdot d(y, A)$$

where K is an upper bound for $d(y, A_\rho)$. This is a formula since we only range over $||y|| \leq \rho + \varepsilon \leq R_y$.

For any $y \in \bar{B}_\rho(A)$, we have

$$\phi_{\rho,\rho+\varepsilon}(x) \leq d(x,y) + 0.$$

So we can say

$$\phi_{\rho,\rho+\varepsilon}(x) \leq d(x, \bar{B}_\rho(A)).$$

We would like to build the required integral using Riemann sums involving $\phi_{\rho,\rho+\varepsilon}$. The bound just established is part of this, but we need to know $\phi_{\rho,\rho+\varepsilon}$ is not too small either. Suppose $\phi_{\rho,\rho+\varepsilon}(x) < d(x, \bar{B}_\rho(A))$. Then we can pick a witness $y \in \bar{B}_{\rho+\varepsilon}(M)$ such that

$$d(x,y) + (K/\varepsilon) \cdot d(y, A) < d(x, \bar{B}_\rho(A)).$$

By definition of $d(y, A)$, for any $\varepsilon' > 0$, there is some $z_{\varepsilon'} \in A$ with

$$|d(y, z_{\varepsilon'}) - d(y, A)| \leq \varepsilon'.$$

For any such $z_{\varepsilon'}$, we have

$$d(x,y) + (K/\varepsilon) \cdot (d(y, z_{\varepsilon'}) - \varepsilon') < K.$$

This implies

$$d(y, z_{\varepsilon'}) < \varepsilon + \varepsilon'.$$

Since $y \in \bar{B}_{\rho+\varepsilon}(M)$, we must have $z_{\varepsilon'} \in \bar{B}_{\rho+2\varepsilon+\varepsilon'}(M)$ and hence

$$z_{\varepsilon'} \in \bar{B}_{\rho+2\varepsilon+\varepsilon'}(A).$$

So, for any $\varepsilon' > 0$, we have $z_{\varepsilon'} \in \bar{B}_{\rho+2\varepsilon+\varepsilon'}(A)$ satisfying the following.

$$\begin{aligned} d(x, z_{\varepsilon'}) &\leq d(x,y) + d(y, z_{\varepsilon'}) \\ &\leq d(x,y) + (d(y,A) + \varepsilon') \\ &\leq d(x,y) + (K/\varepsilon) \cdot d(y,A) + \varepsilon' \cdot (K/\varepsilon) \end{aligned}$$

This gives the following bound for $\inf_{z \in \bar{B}_{\rho+2\varepsilon+\varepsilon'}(A)} d(x,z)$, which is just the distance of x to the closed ball.

$$d(x, \bar{B}_{\rho+2\varepsilon+\varepsilon'}(A)) \leq d(x,y) + (K/\varepsilon) \cdot d(y,A) + \varepsilon' \cdot (K/\varepsilon)$$

Note that for any $\varepsilon^* \geq \varepsilon'$, we can replace the left side like so:

$$d(x, \bar{B}_{\rho+2\varepsilon+\varepsilon^*}(A)) \leq d(x,y) + (K/\varepsilon) \cdot d(y,A) + \varepsilon' \cdot (K/\varepsilon).$$

Since this holds whenever $\varepsilon^* \geq \varepsilon' > 0$, we can fix ε^* and let $\varepsilon' \to 0$, then let $\varepsilon^* \to 0$ to get

$$d(x, \bar{B}_{\rho+2\varepsilon}(A)) \leq d(x,y) + (K/\varepsilon) \cdot d(y,A).$$

This holds for any $y \in \bar{B}_{\rho+\varepsilon}(M)$ witnessing the inf involved in

$$\phi_{\rho,\rho+\varepsilon}(x) < d(x, \bar{B}_\rho(A)).$$

So, if $\phi_{\rho,\rho+\varepsilon}(x) \neq d(x, \bar{B}_\rho(A))$, we have

$$d(x, \bar{B}_{\rho+2\varepsilon}(A)) \leq \phi_{\rho,\rho+\varepsilon}(x) < d(x, \bar{B}_\rho(A)).$$

We have shown that for all $\rho \in [r, r']$ and $\varepsilon > 0$, we have

$$d(x, \bar{B}_{\rho+2\varepsilon}(A)) \leq \phi_{\rho,\rho+\varepsilon}(x) \leq d(x, \bar{B}_\rho(A)).$$

If we fix a bound $||x|| \leq R_x$, then this is sufficient to construct a formula equivalent to

$$\int_r^{r'} d(x, \bar{B}_\rho(A))d\rho$$

when restricted to $||x|| \leq R_x$. For these bounded x, the function $d(x, \bar{B}_\rho(A))$ is bounded and monotonic in ρ, and we can view it as a limit of Riemann sums using the formulas $\phi_{\rho,\rho+\varepsilon}(x)$ and a controlled u. $\square$

Corollary 3.9.4. *If A is a subset of product $\prod M_n$ of sorts, then A is definable if and only if for every finite projection $\pi_N(A)$ and every real N-tuple r, there is an L-formula which interprets as a function equivalent to $d_N((x_1, \dots, x_N), \pi_N(A))$ when restricted to N-tuples $x \in \bar{B}_r(M)$.*

The last couple of propositions suggest the following definition.

Definition 3.9.5. Let ϕ be an L-formula and let M be an L-structure. If the zero set A of the function ϕ^M is definable in M, and we have a collection of formulas $\phi_{r,N}$, each equivalent to $d_N((x_1, \dots, x_N), \pi_N(A))$ when restricted to $x \in \bar{B}_r(M)$, we say that the zero set of ϕ^M is **defined by** the formulas $\phi_{r,N}$.

Let $\mathcal{K}$ be a class of L-structures. We say that the zero set of ϕ is **definable in the class** $\mathcal{K}$ if there are formulas $\phi_{r,N}$ such that for every $M \in \mathcal{K}$, the zero set A_M of the function ϕ^M is defined by the formulas $\phi_{r,N}$.

This next lemma is useful for showing the existence of formulas equivalent to distance functions. In practice, it is often easier to guess a formula with the correct zero set and then use this lemma to obtain the distance function than to explicitly write the distance function at first. We have adapted this lemma from Proposition 9.19 of [3], adjusting its statement and proof for this setting.

Lemma 3.9.6. *Let A be a subset of a finite product $M = \prod_{n=1}^N M_n$ of sorts. Suppose that for a pair of real N-tuples r and r' there is a formula $P_{r,r'}$ such that*

- *$P_{r,r'}(y) = 0$ for all $y \in B_{r'}(A)$, and*
- *for all $\varepsilon > 0$, there is $\delta > 0$ such that for all $y \in B_{r'}(M)$ we have that*

$$P_{r,r'}(y) \leq \delta \text{ implies } d_N(y, B_r(A)) \leq \varepsilon.$$

Then there is an L-formula equivalent to $\int_r^{r'} d_N(x, B_\rho(A))d\rho$ in M.

Proof. Let r and r' be a pair of real N-tuples. By 2.3.12, there is a continuous, increasing function $\alpha : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ such that $\alpha(0) = 0$ and for all $y \in B_{r'}(M)$ we have $d_N(y, B_r(A)) \leq \alpha(P_{r,r'}(y))$. Consider the L-formula F given by

$$F(x) = \inf_y]_r^{r'} \left(\alpha(P_{r,r'}(y)) + d(x,y)\right).$$

By definition of α and the triangle inequality, we have the following.

$$\begin{aligned} F^M(x) &= \int_r^{r'} \inf_{y \in B_\rho(M)} \left(\alpha(P_{r,r'}(y)) + d(x,y)\right) d\rho \\ &\geq \int_r^{r'} \inf_{y \in B_\rho(M)} \left(d(y, B_r(A)) + d(x,y)\right) d\rho \\ &\geq \int_r^{r'} d(x, B_\rho(A)) d\rho \end{aligned}$$

Additionally, since $y \in B_{r'}(A)$ implies $P_{r,r'}(y) = 0$ and hence $\alpha(P_{r,r'}(y)) = 0$, we conclude as follows.

$$\begin{aligned} F^M(x) &= \int_r^{r'} \inf_{y \in B_\rho(M)} \left(\alpha(P_{r'}(y)) + d(x,y)\right) d\rho \\ &\leq \int_r^{r'} \inf_{y \in B_\rho(A)} \left(\alpha(P_{r'}(y)) + d(x,y)\right) d\rho \\ &= \int_r^{r'} \inf_{y \in B_\rho(A)} d(x,y) d\rho \\ &= \int_r^{r'} d(x, B_\rho(A)) d\rho \end{aligned}$$

□

Corollary 3.9.7. *If for all N-tuples r, r' there is $P_{r,r'}$ as above, then A is definable in M.*

When A is bounded, the previous lemma can be applied in the case where r and r' are radii containing the entirety of A. Then the lemma shows the existence of a formula equivalent to $d(x, A)$.

If A is a subset of a countable product of balls, we can extend this.

Corollary 3.9.8. *Suppose A is a subset of a countable product $\prod_{n \in \mathbb{N}} B_{R_n}(M_n)$ of balls in sorts from M. Let*

$$d'(x,y) = \sum_{n \in \mathbb{N}} \frac{d(x_n, y_n)}{1 + d(x_n, y_n)} \cdot 2^{-n}.$$

Suppose there is a formula $P(y)$ such that

- $P(y) = 0$ *for all* $y \in A$

- *For all $\varepsilon > 0$, there is $\delta > 0$ such that for all $y \in M$ we have that*
$$P(y) \leq \delta \text{ implies } d'(y, A) \leq \varepsilon$$
where
$$d'(y, A) = \inf_{a \in A} d'(y, a).$$

Then A is definable.

Proof. Notice that bounds on $d'(x, y)$ imply bounds on the distance in each coordinate. So, we can bound the distance of any finite N-tuple y to $\pi_N(A)$ by knowing $P(y, a) \leq \delta$ for some $a \in \pi_{\mathbb{N} \setminus N}(A)$. This means we can obtain the required formulas using forced limits of $\inf_x] P(y, x)$ formulas. □

The next proposition quickly checks that any finite or countable product definable set is the zero set of some formula. This is a simple consequence of having formulas equivalent to $d(x, A)$ for the finite product case.

Proposition 3.9.9. *Let A be a definable set in M. If A is a subset of a finite product, then A is the zero set of the formula equivalent to $d_N(x, A)$. If A is a subset of a countable product, then each A is the zero set of the formula*
$$\sum_{n=1}^{\infty} \frac{\min(1, d_n(x, \pi_n(A)))}{2^n}$$
where d_n is the max of the distances in the first n coordinates.

Proof. This formula is 0 exactly when every term is 0. This happens exactly when every projection of the sequence x is the projection of an element of A, i.e. $x \in A$. □

The next few propositions will give further useful equivalent characterizations of definable sets. First, we will see that definable sets are the sets we can quantify over as though they were a sort.

Proposition 3.9.10. *Let A be a subset of a finite product $\prod_{n=1}^{N} M_n$ of sorts. The following are equivalent.*

(1) *A is definable in M over C.*
(2) *There is a sequence c over C such that for any formula $\phi(x, y)$ over variable sequences $x = (x_n)_{1 \leq n \leq N}$ and $y = (y_n : n \in \beta)$, and for any real sequences*
$$R = (R_n : n \in \beta),$$
$$r = (r_n)_{1 \leq n \leq N},$$
$$r' = (r'_n)_{1 \leq n \leq N}$$
satisfying $r_n < r'_n$ for all $1 \leq n \leq N$, there is a finite subset $F \subseteq \beta$ and a formula $\psi(y_n, z_n : n \in \mathbb{N})$ such that $\psi^M(y_n, c_n : n \in \mathbb{N})$ is the function
$$\left(\prod_{n=1}^{N} \frac{1}{r'_n - r_n} \right) \int_{r_N}^{r'_N} \cdots \int_{r_1}^{r'_1} \inf_{(x_1, \ldots, x_N) \in \bar{B}_{(\rho_1, \ldots, \rho_N)}(A)} \phi^M(x, y) d\rho_1 \cdots d\rho_N$$

when restricted to y satisfying $||y_n|| \le R_n$ for $n \in F$.
A similar statement holds with sup *in place of* inf.

Proof. ($2 \to 1$) Applying (2) to formulas $d_N(x,y)$ gives the formulas needed in the definition of definability.

($1 \to 2$) We will discuss the single variable x, $\int$ inf case. The finite variable case is an easy adjustment using multivariable inf and $\int$. The $\int$ sup cases are similar.

Let $\phi(x,y)$ be any formula, and let $r < r'$. We are supposing that for all $t < t'$ with $t, t' \in [r, r']$, the function $\int_t^{t'} d(x, \bar{B}_\rho(A))d\rho$ is given by a formula (at least equivalent due to the bounds imposed on x). By uniform continuity of ϕ when restricted to $x \in \bar{B}_{r'}(M)$ and $y \le R$, Lemma 2.3.12 gives us an α such that for any $y \le R$,

$$\phi(x,y) \le \phi(z,y) + \alpha(d(x,z)).$$

Taking $\inf_{x\in\bar{B}_\rho(A)}$ for any $\rho \in [r, r']$, we get

$$\inf_{x\in\bar{B}_\rho(A)} \phi(x,y) \le \phi(z,y) + \alpha(d(z, \bar{B}_\rho(A)))$$

where we have used that α is increasing so that we can pass the inf through. Taking $\inf_{z\in\bar{B}_\rho(M)}$ and comparing it to the larger $\inf_{z\in\bar{B}_\rho(A)}$, we get

$$\inf_{x\in\bar{B}_\rho(A)} \phi(x,y) \le \inf_{z\in\bar{B}_\rho(M)} \left(\phi(z,y) + \alpha(d(z, \bar{B}_\rho(A)))\right) \le \inf_{z\in\bar{B}_\rho(A)} \phi(z,y)$$

where in the right-most term, we have used that $\alpha(0) = 0$. Note that the left-most and right-most terms are equal, so this is an equality. Now, we can integrate over $\rho \in [r, r']$ to get

$$\int_r^{r'} \inf_{x\in\bar{B}_\rho(A)} \phi(x,y)d\rho = \int_r^{r'} \inf_{z\in\bar{B}_\rho(M)} \left(\phi(z,y) + \alpha(d(z, \bar{B}_\rho(A)))\right) d\rho.$$

The left term is the desired function, so it remains to show that the right term is actually a formula.

The issue to address is that $d(z, \bar{B}_\rho(A))$ depends on ρ and moreover when ρ is fixed might not be a formula since it is not the formula guaranteed by (1). But, (1) does give us that

$$\int_t^{t'} d(z, \bar{B}_\rho(A))d\rho$$

is a formula for all $t < t'$.

We will reconstruct the integral on the right side of the above equality by using these $\int_t^{t'} d(z, \bar{B}_\rho(A))$ for values of t, t' in $[r, r']$. To do so, we need to show that we can express the $\int_r^{r'}$ integral in terms of Riemann sums over partitions $\rho_1 < \cdots < \rho_K$ where the averages $\frac{1}{t'-t}\int_t^{t'} d(z, \bar{B}_\rho(A))$ can be used in place of the $d(z, \bar{B}_{\rho_k}(A))$ appearing in terms of the sum.

The main point in this argument is the following. For any $\rho < \rho'$ in $[r, r']$, we can show the following for all $t \in [\rho, \rho']$.

$$\begin{aligned} & \inf_{z \in \bar{B}_{\rho'}(M)} \big(\phi(z,y) + \alpha(d(z, \bar{B}_{\rho'}(A)))\big) \\ \leq \quad & \inf_{z \in \bar{B}_t(M)} \left(\phi(z,y) + \alpha(\frac{1}{\rho' - \rho}\int_\rho^{\rho'} d(z, \bar{B}_\tau(A)))d\tau\right) \\ \leq \quad & \inf_{z \in \bar{B}_\rho(M)} \big(\phi(z,y) + \alpha(d(z, \bar{B}_\rho(A)))\big) \end{aligned}$$

To see this, note that since $d(z, \bar{B}_\tau(A))$ is monotonic in the radius of the ball, the average satisfies

$$d(z, \bar{B}_{\rho'}(A)) \leq \frac{1}{\rho' - \rho}\int_\rho^{\rho'} d(z, \bar{B}_\tau(A))d\tau \leq d(z, \bar{B}_\rho(A)).$$

Also note that α is increasing, that addition of $\phi(z,y)$ preserves the inequalities for fixed z, y, and finally that taking infs also preserves the inequalities.

Now, for every ε, we can find (L, R, ε)-good partitions of $[r, r']$ for

$$\int_r^{r'} \inf_{z \in B_\rho(M)} (\phi(z,y) + \alpha(d(z, B_\rho(A))))$$

as in 2.7.6, depending on the control functions for ϕ, α, the bounds r, r', and bounds for y. We can write the Riemann sums over these partitions and use the above inequalities to replace the $d(z, \bar{B}_\rho(A))$ terms with averages, which are guaranteed to be formulas by (1). This lets us express the integral as a composition of a controlled function with Riemann sums using only formulas, provided we restrict to y with $||y_n|| \leq R_n$ for $n \in F$. Hence,

$$\int_r^{r'} \inf_{x \in \bar{B}_\rho(A)} \phi(x,y)d\rho$$

is given by a formula for bounded y as in the claim. $\square$

An immediate corollary is that when A is bounded, we can take an exact sup and inf over the entirety of A. This is the same trick as for bounded sorts.

Corollary 3.9.11. *If A is a bounded definable subset of a finite product of sorts from M, then for any $\phi(x,y)$, provided bounds on the variables in y, there are formulas equivalent to $\sup_{x \in A} \phi^M(x,y)$ and $\inf_{x \in A} \phi^M(x,y)$.*

Proof. Use the quantifiers from the proposition, with r, r' chosen larger than the magnitudes of coordinates in A. $\square$

The above results justify the introduction of some notation for quantification over definable sets.

Definition 3.9.12. When A is a definable subset of a finite product of sorts in M over C, and ϕ, x, y, R, r, r' are as above, we may denote the formulas $\psi(y_n, z_n : n \in \mathbb{N})$ obtained in the previous proposition by

$$\inf_{x \in A} R]_r^{r'} \phi(x, y)$$

and

$$\sup_{x \in A} R]_r^{r'} \phi(x, y).$$

If A is a bounded subset of a finite product, we write $\sup_{x \in A} \phi$ and $\inf_{x \in A} \phi$ for the formulas from the previous corollary.

This is consistent with the usual quantifier notation with x ranging over sorts. That is, if A is just a product of sorts, then these are equivalent to the usual quantifiers given by the language.

We will finish this section by giving a characterization of definability via preservation of zero sets under ultraproducts. To help state the characterization, we will define ultraproducts of subsets.

Definition 3.9.13. Let $M = \prod_{\mathcal{U}} M_i$, and let $A_i \subseteq M_i$ for all i. We write $\prod_{\mathcal{U}} A_i$ to mean the subset of M consisting of those classes $a \in M$ for which there is some $F \in \mathcal{U}$ and representative $(a_i : i \in I)$ for a having $a_i \in A_i$ for all $i \in F$.

Proposition 3.9.14. *Let ϕ be an L-formula with free variable sequence $x = (x_n : n \in \alpha)$. Let $\mathcal{K}$ be a class of L-structures which is closed under ultraproducts.*

Then the zero set of ϕ is definable in $\mathcal{K}$ if and only if whenever $M = \prod_{\mathcal{U}} M_i$ is an ultraproduct over structures in $\mathcal{K}$ we have

$$\mathrm{Zero}(\phi^M) = \prod_{\mathcal{U}} \mathrm{Zero}(\phi^{M_i})$$

where $\mathrm{Zero}(\phi^N)$ means the zero set of the interpretation of ϕ in the structure N.

Proof. Notice that by the fundamental theorem of ultraproducts, $\prod_{\mathcal{U}} \mathrm{Zero}(\phi^{M_i})$ is always a subset of $\mathrm{Zero}(\phi^M)$, so the statement is only really about the other inclusion.

Suppose the zero set of ϕ is definable in $\mathcal{K}$, let $M = \prod_{\mathcal{U}} M_i$, and suppose $\phi^M(a) = 0$. We assume without loss of generality that ϕ is finitary. Write $a = (a_0, \ldots, a_{N-1})$ and let $(r_0, \ldots, r_{N-1})$ be reals such that $||a_n|| < r_n$ for $n < N$. Let $(a_i : i \in I)$ be a representative sequence for a.

By definability, there is an L-formula $d(x, \mathrm{Zero}(\phi))$ giving the distance from x to the zero set, independently of which structure in $\mathcal{K}$ we use, provided we restrict to x with $||x_n|| < r_n$ for $n < N$. We have

$$d(a, \mathrm{Zero}(\phi^M)) = 0$$

by assumption. So, for any $\varepsilon > 0$ there is $F_{1,\varepsilon} \in \mathcal{U}$ where

$$d(a_i, \mathrm{Zero}(\phi^{M_i})) < \varepsilon \text{ for } i \in F_{1,\varepsilon}.$$

There is also an $F_2 \in \mathcal{U}$ such that

$$||a_{i,n}|| < r_n$$

for all $i \in F_2$ and $n < N$. Thus, for each $\varepsilon > 0$, for $i \in F_{1,\varepsilon} \cap F_2$, the distance from a_i to $\mathrm{Zero}(\phi^{M_i})$ is $< \varepsilon$. This is equivalent to saying that for all $\varepsilon > 0$, for all $i \in F_{1,\varepsilon} \cap F_2$, there is some

$$b_i^{(\varepsilon)} \in \mathrm{Zero}(\phi^{M_i})$$

with

$$d(a_i, b_i) < \varepsilon.$$

This shows that we can inductively construct another representative sequence $(b_i : i \in I)$ for the class a, with $b_i \in \mathrm{Zero}(\phi^{M_i})$ holding on some $F \in \mathcal{U}$. This demonstrates the needed inclusion when the zero set is definable.

Suppose the zero set of ϕ is not definable in $\mathcal{K}$. If ϕ itself satisfied the ε, δ property in Lemma 3.9.6 independently of the structure $M \in \mathcal{K}$, then $\mathrm{Zero}(\phi)$ would be definable in $\mathcal{K}$, so this must not be the case.

Instead, it must be possible to find some positive real tuple $r = (r_1, \ldots, r_N)$ and some $\varepsilon > 0$ such that for all $\delta > 0$, there is $M_\delta \in \mathcal{K}$ with $b_\delta \in \bar{B}_{(r_1,\ldots,r_N)}(M_\delta)$ having

$$\phi(b) \leq \delta$$

but

$$d(b, \mathrm{Zero}(\phi^{M_\delta})) > \varepsilon.$$

Consider a nonprincipal ultraproduct $M = \prod_{\mathcal{U}} M_{1/n}$ with $\mathcal{U}$ over $\mathbb{N}$ formed using such M_δ. This yields a structure in $\mathcal{K}$ containing a point $b = (b_{\frac{1}{n}} : n \in \mathbb{N})/\sim$. By the fundamental theorem of ultraproducts, this b satisfies

$$b \in \bar{B}_{(r_1,\ldots,r_N)}(M)$$

and

$$\phi^M(b) = 0$$

but also

$$b \notin \prod_{\mathcal{U}} \mathrm{Zero}(\phi^{M_{1/n}}).$$

This shows that the inclusion fails when the zero set is not definable. $\square$

3.10. **Extension by constants.** We often want to have constant symbols in our language, or are interested in the behavior of formulas $\phi(x, y)$ where we allow x to vary but keep y fixed. We discuss these situations in this section.

To add a constant to a language, it suffices to add a 0-ary function symbol c. It is important to remember that function symbols must have controllers known to the language. The controllers N and δ are trivial since c is 0-ary, but notably, the scaling modulus λ amounts to a bound for the magnitude of the intended constant c.

The following definition helps when we discuss formulas with fixed parameters.

Definition 3.10.1. Let L be a language, and let C be a set with a corresponding sort of L and bound in $\mathbb{R}_{\geq 0}$ for $||c||$ assigned to each $c \in C$. We define $L(C)$ to be the language obtained by extending L so that it has a constant symbol for each $c \in C$. We call $L(C)$ the **extension of** L **by constants** C.

If M is an L-structure and $C \subseteq M$, we write $M(C)$ to denote an $L(C)$-structure obtained by extending M to interpret each symbol c as the corresponding constant element c in M. In this case, we define the controller λ to be the function with constant value $||c||$.

We write $\mathrm{Th}(M(C))$ to denote the $L(C)$-theory of $M(C)$.

Note that since the theory has been taken from a specific model, $\mathrm{Th}(M(C))$ will contain the sentence $d(c, \star) \approx r$ where r is the real number $d^M(c^M, \star^M)$. So, in any other model of this theory, c must have the same distance from the basepoint.

3.11. **Types and saturation.** If a is in an L-structure M, then L can "talk about" a and how it relates to the rest of M using formulas and possibly parameters from M. That is, $L(C)$ provides some kind of description for a given by the values of $\phi^M(a)$ for each $L(C)$-formula ϕ.

Conversely, we might have a collection p of $L(C)$-formulas and corresponding values, and wonder whether there is any a which is described by p in this sense. Treating these collections p as objects themselves is a useful possibility because of our use of languages.

Readers unfamiliar with logic are likely familiar with the Dedekind cut construction of the reals from the rationals and might appreciate this analogy. A cut can be thought of as a description of some kind of ideal point using $<$ and rational parameters. Every rational q has a corresponding cut, but not all cuts correspond to a rational.

Below, we will use notation as though the language has only one sort, but remember that the language could be multi-sorted, and variables can have different sorts.

Definition 3.11.1. Let L be a language, and let $L(C)$ be some extension of L by constants. For any $L(C)$-structure M and (finite or infinite) tuple a over

M, we define

$$\mathrm{tp}_M(a/C)$$

to be the set of $L(C)$-formulas $\phi(x)$ such that $M \models \phi(a)$.

We say p is a **type** over C if there is some $L(C)$-structure M and tuple a over M such that $p \subseteq \mathrm{tp}_M(a/C)$. If $p = \mathrm{tp}_M(a/C)$, we say p is a **complete type**. In either case, we say a **realizes** p in M.

If a is an n-tuple for some $n \in \mathbb{N}$, or equivalently if p consists of formulas involving only n-many free variables $x_1, x_2, \ldots, x_n$, then we call p an n-**type**. More generally, if we use a collection of variables indexed by a set I, we call p an I-**type**.

We are usually interested in types which are realized in structures restricted to a certain class. For example, if we are discussing Hilbert spaces, we might only care about types that can be realized in a Hilbert space.

Definition 3.11.2. If we have a class $\mathcal{T}$ of $L(C)$-structures, we say p is a **type** in $\mathcal{T}$ over C if there is an $L(C)$-structure $M \in \mathcal{T}$ such that $p \subseteq \mathrm{tp}_M(a_1, a_2, \ldots/C)$.

We denote the set of complete n-types in $\mathcal{T}$ over C by $S_n(\mathcal{T})$. Similarly for I-types.

Often, $\mathcal{T}$ is the set of models of some theory $\mathrm{Th}(M(C))$. In this case, we just write $S_n(C)$ when the model and language are clear.

To try to summarize things in an intuitive way, types in $S_n(C)$ are the possible, complete descriptions of n-tuples that exist in at least one structure elementarily equivalent to $M(C)$.

Next, we define saturation. Informally, saturated structures have a kind of logical completeness property. We will see after this definition that it suffices to only consider 1-types; we get I-types for free.

Definition 3.11.3. Let M be an L-structure, and let κ be an infinite cardinal. We say M is κ-**saturated** if for any index set I having cardinality $\leq \kappa$ and subset $C \subseteq M$ having density $< \kappa$, every type in $S_I(\mathrm{Th}(M(C))$ is realized in M.

Proposition 3.11.4. *M is κ-saturated iff for any $C \subseteq M$ with density $< \kappa$, all types in $S_1(C)$ are realized in M.*

Proof. The $(\rightarrow)$ direction is trivial, so we just check $(\leftarrow)$. The proof is by transfinite induction on I. The idea is to realize the tuple one variable at a time, pushing the partial realizations into the constants of the language.

The base case is given by the assumption. Suppose we know all types in $S_\alpha(C)$ are realized for some $\alpha < \kappa$. Let $p \in S_{\alpha+1}(C)$, and we will show that it is also realized. Notice that p ranges over variables x_β with $\beta < \alpha + 1$. Take $q \in S_\alpha(C)$ to be the subset of p given by

$$q = \{\phi \in p : \phi \text{ uses only variables } x_\beta \text{ with } \beta < \alpha\}.$$

That is, q refers to all variables except x_α, and q is clearly a type since it is realized by the truncation of whatever realizes p. By the inductive hypothesis, we have some $(a_\beta : \beta < \alpha)$ realizing q in M. Now, consider the set r given by

$$r = \{\phi(x_\alpha, a_\beta : \beta < \alpha) : \phi \in p\}.$$

That is, r is the set of formulas from p with a_β substituted for x_β when $\beta < \alpha$. Since p consisted of $L(C)$-formuals, r consists of $L(C \cup \{a_\beta : \beta < \alpha\})$-formulas. But $C \cup \{a_\beta : \beta < \alpha\}$ is still of cardinality $< \kappa$. We will check that r is a type, so that $r \in S_1(C \cup \{a_\beta : \beta < \alpha\})$ and hence must be realized by M. To see this, just notice that since $p \in S_{\alpha+1}(C)$, there is some $N \models \mathrm{Th}(M(C))$ with $(b_\beta : \beta < \alpha + 1)$ realizing p in N. Since we have

$$\begin{aligned} q &= \mathrm{tp}_M(a_\beta : \beta < \alpha) \\ &\subseteq \mathrm{tp}_N(b_\beta : \beta < \alpha + 1) \\ &= p \end{aligned}$$

we can extend N to an $L(C \cup \{a_\beta : \beta < \alpha\})$-structure modeling

$$\mathrm{Th}(M(C \cup \{a_\beta : \beta < \alpha\}))$$

by interpreting the symbol for a_β as the element b_β. Then r is realized in this extension of N by the element b_α from our original tuple $(b_\beta : \beta < \alpha + 1)$. So, we have that r is realized by some a_α in M. Finally, just observe that $(a_\beta : \beta < \alpha + 1)$ realizes p in M due to the construction of r and q from p.

Say $p \in S_\alpha(C)$ for some limit ordinal $\alpha \le \kappa$. Then

$$p = \mathrm{tp}_N(b_\beta : \beta < \alpha)$$

for some $N \models \mathrm{Th}(M(C))$. A similar argument to the above shows we can construct the required tuple in M by starting with a partial realization and then extending it. For example, if $p \in S_\omega(C)$, we can take a finite tuple in M realizing the restriction of p to the first n-variables for some n, then inductively realizing extensions to more variables by putting the partial realization into the constants of our language. This generates an sequence realizing p in M. $\square$

3.12. **Homogeneity.** Say M_C is an $L(C)$-structure, with M the structure obtained by ignoring the constants in C, and suppose $a, b \in M$. If σ is an $L(C)$-automorphism of M_C with

$$\sigma(a) = b$$

then we will have

$$\mathrm{tp}_{M_C}(a) = \mathrm{tp}_{M_C}(b)$$

because $L(C)$-isomorphisms must preserve all $L(C)$-formulas. Phrased differently but equivalently, we are observing that if an L-automorphism σ fixes some subset $C \subseteq M$ pointwise but sends a to b, then

$$\mathrm{tp}_M(a/C) = \mathrm{tp}_M(b/C).$$

This section addresses some definitions and propositions regarding the converse of this situation. That is, if $\mathrm{tp}_M(a/C) = \mathrm{tp}_M(b/C)$, is there a C-fixing L-automorphism of M which sends a to b? The question is essentially whether the partial map $a \mapsto b$ extends to an automorphism.

Definition 3.12.1. Let M be an L-structure and let κ be an infinite cardinal. We call M **strongly κ-homogeneous** if whenever $L(C)$ is an extension of L by constants with $\mathrm{card}(C) < \kappa$ and whenever M_{C_1}, M_{C_2} are two $L(C)$-structures obtained by adding interpretations to M for the constants in C, we have that

$$\begin{array}{ccc} \text{elementary equivalence} & & \text{isomorphism} \\ M_{C_1} \equiv M_{C_2} & \text{implies} & M_{C_1} \cong M_{C_2}. \end{array}$$

Because M_{C_1} and M_{C_2} are both extensions by constants of the same structure M, this definition can also be read as a statement about elements a, b in M sharing the same type over some parameters $C \subseteq M$. If we consider the $L(C \cup \{c^*\})$-structures $M_{C,a}$ and $M_{C,b}$ obtained by interpreting each constant in C as the corresponding element of M, but intepreting c^* as a in $M_{C,a}$ and as b in $M_{C,b}$, then we have $M_{C,a} \equiv M_{C,b}$ whenever $\mathrm{tp}_M(a/C) = \mathrm{tp}_M(b/C)$.

The following proposition guarantees that every L-structure has elementary extensions which are κ-saturated and strongly κ-homogeneous. This should be seen as the analog in model theory of the closure or completeness results in other areas of mathematics. This result actually subsumes many of those when the language and structures are appropriately chosen.

Proposition 3.12.2. *Let M be an L-structure and let κ be an infinite cardinal. Then M has a κ-saturated, strongly κ-homogeneous elementary extension N.*

Proof. See Proposition 7.12 of [3]. The proof is the same in this setting. □

3.13. **Type spaces.** Remember that $S_n(\mathcal{T})$ is the set of complete n-types which can be realized in some structure from $\mathcal{T}$. Here we will look at the logic topology and d-metric on $S_n(\mathcal{T})$.

Definition 3.13.1. Let $\mathcal{T}$ be a class of $L(C)$-structures, and fix some $n \in \mathbb{N}_+$. For each nonnegative $L(C)$-formula ϕ and $\varepsilon > 0$, we define

$$[\phi < \varepsilon]$$

to be the set

$$\{p \in S_n(\mathcal{T}) : (\phi \lesssim \delta) \in p \text{ for some } \delta < \varepsilon\}.$$

That is, p is in $[\phi < \varepsilon]$ exactly when all realizations $a \in M$ of p must have $\phi(a)^M < \varepsilon$. The **logic topology** on $S_n(\mathcal{T})$ is the topology with open sets generated by the $[\phi < \varepsilon]$.

Remark 3.13.2. Note that we are only using nonnegative ϕ in the above definition. We could obtain the same topology by instead allowing all formulas but only considering the sets

$$[|\phi| < \varepsilon].$$

Convergence of a sequence of types $p_1, p_2, \dots$ to a type q in the logic topology requires that for any $\phi \in q$ and $\delta > 0$, eventually $(\phi \lesssim \delta) \in p_n$. That is, (p_n) must eventually agree arbitrarily well with q on any formula.

Now we will discuss the d-metric for type spaces. This is a metric on types determined by the "ideal" distance between their realizations.

Definition 3.13.3. Let T be an $L(C)$ theory. The d-**metric** on $S_n(T)$ is defined by setting

$$d(p,q) = \min \left\{ d^M(a,b) \text{ such that } \left(\begin{array}{c} M \models T \\ a \text{ realizes } p \text{ in } M \\ b \text{ realizes } q \text{ in } M \end{array} \right) \right\}$$

where as usual d^M is taken to be the max metric on coordinates for n-tuples. The **basepoint** of $S_n(T)$ is defined to be the type of the basepoint sequence, i.e.

$$\mathrm{tp}((\star, \star, \dots, \star)/C).$$

Equivalently, $d(p,q)$ is the minimal r such that the formula

$$\max\{d(x_1,y_1), \dots, d(x_n,y_n)\} \approx r$$

appears in a common extension of p, q to a $2n$-type. We can say "minimal" rather than "infimal" in the definition because an infimal r could be witnessed exactly, say by taking ultraproducts.

Proposition 3.13.4. *Let T be an $L(C)$-theory. The d-metric is complete on $S_n(T)$ and induces a finer topology than the logic topology.*

Proof. Suppose p_k is a cauchy sequence of types over C in $S_n(T)$, and without loss of generality, we suppose further that $d(p_k, p_{k+1}) \leq 2^{-k}$ for all k. Let M be an ω-saturated and strongly ω-homogeneous model of T. Note that M is already an $L(C)$-structure, i.e. C is part of the language, so ω is large enough to realize each p_k. Then for each k, by saturation there are realizations a_k, a_{k+1} in M of p_k, p_{k+1} respectively with $d(a_k, a_{k+1}) \leq 2^{-k}$.

Now, M is complete, so if a_k were a cauchy sequence, we could take the type of its limit to finish the proof. However, in realizing these pairs we might have chosen a_k, a_{k+1} and a'_{k+1}, a_{k+2} with $a_{k+1} \neq a'_{k+1}$ and, for example, we are not immediately able to bound $d(a_k, a_{k+2})$. We will inductively construct a sequence a_k applying homogeneity along the way to remedy this problem.

Start with $a_0, a_1 \in M$ realizing p_0, p_1 with $d(a_0, a_1) \leq 2^{-0}$. Given $a_0, \dots, a_k$, choose a'_k, a'_{k+1} realizing p_k, p_{k+1} with $d(a'_k, a'_{k+1}) \leq 2^{-k}$. By homogeneity, there is an automorphism fixing $a_0, \dots, a_{k-1}$, sending a'_k to a_k, and sending a'_{k+1} to some element a_{k+1} which must still realize p_{k+1} and have $d(a_k, a_{k+1}) \leq 2^{-k}$. This inductively constructed sequence a_k is thus a cauchy sequence in M, which converges to some limit a whose type is the limit in $S_n(T)$ of the sequence p_k.

Now we will check that the d-metric has a finer topology than the logic topology. Let $[\phi < \varepsilon]$ be a basic open set in the logic topology, and let $p \in$

$[\phi < \varepsilon]$. In particular, this tells us that p contains some $\phi \lesssim \delta$ with $\delta < \varepsilon$. Since ϕ is controlled and p contains some $d(x,\star) \leq R$, we can use the uniform continuity of ϕ in balls of radius R to determine some positive $\delta' < 1$ such that $||x||, ||y|| < R+1$ and $d(x,y) \leq \delta'$ implies $|\phi(x) - \phi(y)| \leq \frac{\varepsilon-\delta}{2}$. For this δ', we thus have that $d(p,q) \leq \delta'$ implies that q contains $\phi \lesssim \delta + \frac{\varepsilon-\delta}{2}$, and hence $q \in [\phi < \varepsilon]$. We have shown that each basic open set in the logic topology contains an open d-metric ball around every point and so must be an open set in the d-metric topology. □

As we are about to see, for theories with a bound on the diameter of the structure, the type spaces are compact in the logic topology. In general, when our theory includes unbounded structures, we only obtain compactness of the type spaces when restricted to closed d-metric balls around $\mathrm{tp}(\star/C)$. For example, consider the theory of unbounded structures and the open cover

$$\{[d(x,\star) < r] : r > 0\}$$

of the type space S_1. Any finite subset of this cover clearly misses types of points with large $d(x,\star)$. This makes the following proposition a bit awkward to state, since the balls referenced are determined by the d-metric, but the compactness is with respect to the logic topology. However, the reader should keep in mind that complete types always determine their magnitude since they must contain a formula of the form

$$d(x,\star) \approx r.$$

So it is natural to consider types which are bounded in the d-metric.

Proposition 3.13.5. ***(Compactness in the type spaces)*** *Let T be an $L(C)$ theory, let $n \geq 1$, and let $\bar{B}$ be a closed d-metric ball around the basepoint in $S_n(T)$.*

Then $\bar{B}$ is compact and hausdorff in the logic topology on $S_n(T)$.

In particular, if there is some $D \in \mathbb{R}$ such that all models of T have diameter $\leq D$, then $S_n(T)$ is compact and hausdorff in the logic topology.

Proof. This will follow from the ultraproduct theorem or compactness theorem from earlier. For expositional purposes, we will explain both approaches. We will write $\bar{B}(M)$ for the points in a model M realizing at least one type in $\bar{B}$. If we have a cover of $\bar{B}$ by base open sets $[\phi_i < \varepsilon_i]$ for $i \in I$, then every type $p \in \bar{B}$ is in some $[\phi_i < \varepsilon_i]$. It follows that no $p \in \bar{B}$ contains $\phi_i \gtrsim \varepsilon_i$ for all i. This means that

$$P = \{\phi_i \gtrsim \varepsilon_i : i \in I\}$$

is not satisfiable in any model M by an element in $\bar{B}(M)$.

First, the ultraproduct approach. If for every finite subset F of P there were a model M_F with $a_F \in \bar{B}(M)$ realizing F, we could construct an ultraproduct and obtain a point

$$((a_F)/\sim) \text{ in the ball } \bar{B}(M^{\mathcal{U}})$$

which would realize P by the fundamental theorem of ultraproducts. The type

$$\mathrm{tp}_{M^{\mathcal{U}}}\left((a_F)/\sim\right)$$

is then a type in $\bar{B}$ containing P, which is a contradiction. So, there must be a finite subset of P which is not realizable in any model M by an element in $\bar{B}(M)$. Equivalently, there is a finite subset

$$\{\phi_j \gtrsim \varepsilon_j : j = 1, \dots, k\}$$

of P such that no $p \in \bar{B}$ contains $\phi_j \gtrsim \varepsilon_j$ for $j = 1, \dots, k$. Thus, every $p \in \bar{B}$ satisfies

$$p \in \bigcup_{j=1}^{k} [\phi_j < \varepsilon_j].$$

This verifies the compactness of $\bar{B}$.

Now, the alternative approach using the compactness theorem. Let $L(Ca)$ be the language obtained by adding a constant a which is bounded by the radius of $\bar{B}$. Let T_a be T but viewed as a set of $L(Ca)$ sentences. Since P is not satisfiable in any M by an element in $\bar{B}(M)$, it follows that P is not satisfiable among models of T_a when viewed as a collection of sentences in the language $L(Ca)$. By the compactness theorem, P is not finitely satisfiable in T_a. One can now check that P is not finitely satisfiable in models of T and finish as in the ultraproduct case.

The type space being hausdorff can be seen by noticing that distinct types must have at least one ϕ on which they disagree by a positive amount. That is, $\phi \in p$ and $\phi \notin q$ means the there is an $r \in \mathbb{R}$ such that the formula $\phi \approx 0$ is in p but the formula $\phi \approx r$ is in q for some $r \in \mathbb{R}$. Then the open sets

$$[(\phi \approx 0) < \frac{r}{2}]$$

and

$$[(\phi \approx r) < \frac{r}{2}]$$

separate p and q. □

Note that in general, closed d-metric balls in $S_n(T)$ are not compact with respect to the d-metric. This can be seen by considering an infinite structure with the discrete metric. A sequence of types, even when confined to a closed ball, does not necessarily accumulate in the d-metric.

3.14. **Conservative Extensions.** The purpose of viewing mathematical objects as a class of L-structures or models of some L-theory T is to focus on some particular behavior of those objects captured by L, use model theory to find other nicer structures to work with, then use the language and theory as a pivot to transfer our argument to the original object.

The choice of L and T is a trade off: we need to know we have captured the right objects, but we want to keep L and T somewhat low in complexity. This typically means we add just enough symbols to L to get what we need. This makes L easier to understand, and so for example, L-isomorphisms are easier to understand. But it may mean that the interesting behavior of our structures is captured in a complicated way, buried in deeply quantified, complex formulas.

This can result in devoting the initial discussion of models of T to understanding special subsets, relations, and functions that are implicitly captured by L, even though L did not explicitly have sorts or symbols for them. In practice, we end up justifying acting like L has these extra sorts and symbols so that we do not have to keep writing long formulas. The justification for this is nontrivial. In general, if one arbitrarily extends L or T, then the resulting structures and models are more restrictive.

In this section, we discuss a formalization of this extension phenomenon. This is a more category theoretic view of models of T, so we will start with some category theory definitions. This material can be found starting around section 3.8 in Jean-Martin Albert's PhD thesis [1] or Bradd Hart's lecture slides for his class Math 712, which will hopefully remain available online for the duration of the internet.

Definition 3.14.1. Let C and D be categories, and let $F : C \to D$ be a functor.

(1) F is **full** if for all $c_1, c_2 \in C$, the map induced by F between hom sets
$$F : \mathrm{Hom}(c_1, c_2) \to \mathrm{Hom}(F(c_1), F(c_2))$$
is onto.

(2) F is **faithful** if for all $c_1, c_2 \in C$, the map
$$F : \mathrm{Hom}(c_1, c_2) \to \mathrm{Hom}(F(c_1), F(c_2))$$
is one-to-one.

(3) F is **dense** if for all $d \in D$, there is $c \in C$ such that $F(c) \cong d$.

(4) F is an **equivalence of categories** if F is full, faithful, and dense. In this case, we say that C and D are **equivalent categories**.

We will be interested in functors going from the category of models of T' to the category of models of T, where T' is some extension of T to a larger language.

Definition 3.14.2. Let L' be a language extending L, let T' be an L'-theory, let T be an L-theory, and suppose T' contains T. Define $\mathrm{Mod}(T')$ to be the category of models of T' with morphisms given by elementary maps. Similarly, define $\mathrm{Mod}(T)$ to be the category of models of T with morphisms given by elementary maps. The **forgetful functor** $F : \mathrm{Mod}(T') \to \mathrm{Mod}(T)$ is the functor which

- takes any model M' of T' and sends it to the model M of T obtained by ignoring the interpretations of any new symbols in L' over L, and
- takes any elementary map $\mathcal{E}' : M' \to N'$ and sends it to the map $\mathcal{E} : F(M') \to F(N')$ obtained by restricting $\mathcal{E}'$ to the L-sorts of $F(M')$.

That the obtained M actually models T and that F is actually a functor are easy to see from the containments $L \subseteq L'$ and $T \subseteq T'$. In this context, because categories of models have elementary maps as the morphisms and are closed under isomorphism, F is an equivalence of categories if the following things are true.

- (F is full) For any $M, N \models T$ and choice of extensions $M', N' \models T'$ with $F(M') = M$ and $F(N') = N$, any elementary map
$$\mathcal{E} : M \to N$$
lifts to an elementary map
$$\mathcal{E}' : M' \to N'$$
with $F(\mathcal{E}') = \mathcal{E}$.
- (F is faithful) Such liftings of elementary maps are unique.
- (F is dense) Every $M \models T$ lifts to some $M' \models T'$ with $F(M') = M$.

Note that it is possible for there to be distinct extensions of $M \models T$ to models $M', N' \models T'$, but these extensions must be isomorphic if F is an equivalence of categories. This is because the identity map id_M must lift to maps $f : M' \to N'$ and $g : N' \to M'$, and if $g \circ f \neq \mathrm{id}_{M'}$, then there are nonunique liftings of id_M to maps $M' \to M'$.

The topic of this section is when the extension of L to L' and of T to T' does not change the category of models in a meaningful way. That is, we are interested in when $\mathrm{Mod}(T')$ and $\mathrm{Mod}(T)$ are equivalent categories, despite L' possibly having additional sorts, functions, and relations.

For the definition below, recall that $L'(M')$ and $L(M)$ denote extension of a language by constants, in this case extension of L' by constants for each element of M', and extension of L by constants for each element of M.

Definition 3.14.3. If $M \models T$ is the reduct of $M' \models T'$ from L' to L, then we say M is **stably embedded** in M' if for every $L'(M')$-formula $\phi(x)$ with variable tuple x ranging over L-sorts, and for every $\varepsilon > 0$, there is an $L(M)$-formula $\psi(x)$ with the same variables, such that $\phi^{M'}(x) = \psi^{M'}(x)$.

Being stably embedded means that any formula in the extended setting is "already there" in the restricted setting. For example, if L' has an extra relation symbol R, stable embeddedness implies that there is an L-formula ψ which is equivalent to R. Moreover, even if ϕ references elements in M' from L'-sorts, stable embeddedness implies that there are elements from L-sorts of M which suffice. The following proposition tells us that we get stable embeddedness of all reducts if the forgetful functor is full and faithful.

Proposition 3.14.4. *Let $L \subseteq L'$ be languages, $T \subseteq T'$ be theories, and let F be the forgetful functor from models of T' to models of T. If F is full and faithful, then for all $M' \models T'$, the reduct $F(M')$ is stably embedded in M'.*

Proof. Lemma 3.8.2 of [1]. □

Finally, we have the theorem which tells us that if F is an equivalence of categories, then all of the extra structure in $M' \models T'$ is in some sense "already there". For a fuller discussion of interpretability and T^{eq} in the continuous setting, see [1].

Proposition 3.14.5. (Conceptual completeness) Let $L \subseteq L'$ be languages and $T \subseteq T'$ be L- and $L'-$ theories, respectively. Suppose the forgetful functor $F : \mathrm{Mod}(T') \to \mathrm{Mod}(T)$ is an equivalence of categories. Then T' is interpretable in T^{eq}.

Proof. Theorem 3.8.3 of [1]. □

3.15. sup **in a theory.** This section collects some results that make writing axioms for a theory more convenient. Some $\forall$-style quantification is possible using sup] quantifiers. The main idea is that if the average sup of a nonnegative formula $\phi(x)$ is 0 on any bounded part of the space, then it must be that $\phi(x)$ is exactly 0 everywhere.

Proposition 3.15.1. *Fix some L-formula $\psi(x)$, and for each n, let ϕ_n be the L-sentence*

$$\sup_x]_n^{n+1}\psi(x).$$

Then for any $C \in \mathbb{R}$ and any L-structure M, we have

$$\phi_n^M \le C \text{ for all } n$$

if and only if

$$\sup_{x\in M} \psi(x)^M \le C.$$

Proof. If any ϕ_n^M were $> C$, then considering the integral we must have some $x \in M$ with $\psi^M(x) > C$. Conversely, if $\sup_{x\in M} \psi^M > C$, then because $\sup_{x\in B_r(M)} \psi^M$ is a nondecreasing function of r, any integral over $[n, n+1]$ with $x \in B_n(M)$ will be $> C$. □

Corollary 3.15.2. *With the above notation, if $\psi^M(x) \ge 0$ for all $x \in M$, then the following are equivalent.*

- $\phi_n^M = 0$ *for all n.*
- $\sup_{x\in M} \psi^M(x) = 0$.
- $\phi_n \in \mathrm{Th}(M)$ *for all n.*
- *for all $x \in M$, $\psi^M(x) = 0$.*

The above method lets us mimick a $\forall$ quantifier over the entirety of our space. We might want to try mimicking $\forall$ over a bounded part of the space. As usual, there are issues with quantifying over balls. However, we can always find axioms giving the following version of universal quantification over open balls.

Proposition 3.15.3. *Let $R > 0$, and for each $n \in \mathbb{N}_+$, let ϕ_n be the L-sentence*

$$\sup_x]_{R-\frac{1}{n}}^{R} \psi(x)$$

for some fixed $\psi(x)$.

Then for any $C \in \mathbb{R}$ and any L-structure M, we have

$$\phi_n^M \leq C \text{ for all } n$$

if and only if

$$\sup_{x \in B_R(M)} \psi(x)^M \leq C.$$

Proof. Suppose $\phi_n^M > C$ for some n. Considering the integral, there must be some $t \in [R - \frac{1}{n}, R]$ where $\sup_{x \in B_t(S)} \psi(x) > C$. Conversely, suppose $\sup_{x \in B_R(M)} \psi(x)^M > C$. Let $a \in B_R(M)$ witness this sup, so that $\psi(a)^M > C$. Since the sup is nondecreasing as a function of the radius, it follows that whenever $R - \frac{1}{n} > |a|$, we have

$$n \cdot \int_{R-\frac{1}{n}}^{R} \sup_{x \in B_\rho(M)} \psi(x)^M d\rho$$

bounded below by

$$n \cdot \int_{R-\frac{1}{n}}^{R} \psi(a)^M$$

which is $> C$. □

Corollary 3.15.4. *With the above notation, if $\psi^M(x) \geq 0$ for all $x \in M$, then the following are equivalent.*

- $\phi_n^M = 0$ *for all n.*
- $\sup_{x \in B(M)} \psi(x)^M = 0$.
- $\phi_n \in \mathrm{Th}(M)$ *for all n.*
- *For all $x \in B_R(M)$, we have $\psi^M(x) = 0$.*

3.16. inf **in a theory.** This section is like the last but handles inf and existence. We only provide an approximate version of existential quantification on a given bounded part of the space. Recall again that an ultraproduct example can be used to demonstrate problems inherent in existential quantification over the entire space.

Proposition 3.16.1. *Let $R > 0$, and for $n \in \mathbb{N}_+$ let ϕ_n be the L-sentence*

$$\inf_x]_{R}^{R+\frac{1}{n}} \psi(x)$$

for some fixed $\psi(x)$.

Then for any $C \in \mathbb{R}$ and any L-structure M, we have

$$\phi_n^M \leq C \text{ for all } n$$

if and only if

$$\inf_{x \in B_{R+\varepsilon}(M)} \psi^M \leq C \text{ for all } \varepsilon > 0.$$

Proof. Suppose $\phi_n^M > C$ for some n. Then considering the integral there must be some $t \in [R, R + \frac{1}{n}]$ with $\inf_{x \in B_t(M)} \psi^M > C$. Rewriting this, we see that $\inf_{x \in B_{R+(t-R)}(M)} \psi^M > C$, so taking $\varepsilon = t - R$ establishes this direction of the proof. Conversely, suppose $\inf_{x \in B_{R+\varepsilon}(M)} \psi^M > C$ for some $\varepsilon > 0$. Since the inf is nonincreasing as a function of the radius, it follows that whenever $r < \varepsilon$, we have $\inf_{x \in B_{R+r}(M)} \psi^M > C$. So, provided $\frac{1}{n} < \varepsilon$, we have $\inf_{x \in B_r(M)} \psi^M > C$ for all $r \in [R, R + \frac{1}{n}]$. This implies $\phi_n^M > C$. ☐

Corollary 3.16.2. *With the above notation, if $\psi^M(x) \geq 0$ for all $x \in M$, then the following are equivalent.*

- $\phi_n^M = 0$ *for all n.*
- $\inf_{x \in B_{R+\varepsilon}(M)} \psi^M = 0$ *for all $\varepsilon > 0$.*
- $\phi_n \in \mathrm{Th}(M)$ *for all n.*
- *For all $\varepsilon, \varepsilon' > 0$, there exists $x \in B_{R+\varepsilon}(M)$ with $\psi^M(x) \leq \varepsilon'$.*

To rephrase the last statement informally: there are arbitrarily good points arbitrarily close to $\bar{B}_R(M)$.

As usual, if the structure is bounded, things are a bit simpler.

Corollary 3.16.3. *Let M be an L-structure whose underlying metric space is bounded with diameter $\leq D$. Let ϕ be an L-sentence of the form*

$$\inf_{x \in B_r(S)}]_D^{D+1} \psi$$

where ψ is an L-formula with $\psi^M(x) \geq 0$ for all $x \in M$.

Then the following are equivalent.

- ϕ *is in* $\mathrm{Th}(M)$.
- $\inf_{x \in M} \psi = 0$.
- *For all $\varepsilon > 0$, there exists $x \in M$ with $\psi^M \leq \varepsilon$.*

Proof. Just note that $\inf_{x \in B_r(M)} \psi^M$ is equivalent to $\inf_{x \in M} \psi^M$ for $r \geq D$ and apply the above corollary. ☐

4. Stability

Stability is a way of quantifying the complexity of a theory, and it does so by considering the size of type spaces over parameter sets of various sizes. A standard example in the discrete case is that $(\mathbb{Q}, <)$ is unstable, because there is a set of $\aleph_0$ many parameters (the rationals) over which we can find a type space of size $2^{\aleph_0}$ (one for each cut).

In the continuous setting the important notion is not the cardinality of the type space but rather the density of the type space. The notion of density requires a choice of topology, but we have given two different natural choices for the topology of type spaces. We will use the d-metric topology, because as observed in [3], this leads to the results one would expect when generalizing from the discrete setting.

Definition 4.0.1. Let T be an L-theory, and let λ be an infinite cardinal. We say T is **λ-stable** if for any model $M \models T$ and subset $A \subseteq M$ with cardinality $\leq \lambda$, the density of $S_1(T(A))$ with respect to the d-metric is $\leq \lambda$.

We say T is **stable** if it is λ-stable for some λ. If T is not stable we say it is **unstable**.

Stability theory is a rich area of model theory, and there are many ways to approach the subject. We will be primarily interested in a characterization of stability via the order property.

Definition 4.0.2. Let T be an L-theory. We say T has **the order property** if there is

- an L-formula $\phi(x, y)$ with finite variable tuples x, y of the same length,
- a model $M \models T$, and
- a bounded sequence $(a_n)_{n \in \mathbb{N}}$ of tuples, with each a_n of the same length as x and y

such that

$$\phi(a_i, a_j) = \begin{cases} 0 & \text{whenever } i < j \\ 1 & \text{whenever } i \geq j \end{cases}.$$

By a bounded sequence of tuples, we mean there is some real tuple R such that the coordinates of each tuple a_n are always bounded by the corresponding real in R.

Since formulas are controlled, for any such sequence above there must be a minimal distance $\varepsilon > 0$ between each pair a_i, a_j, that is, $d(a_i, a_j) \geq \varepsilon$ whenever $i \neq j$.

Definition 4.0.3. We will say a sequence $(a_n)_{n \in \mathbb{N}}$ is **ε-separated** if

$$d(a_i, a_j) \geq \varepsilon$$

for all $i \neq j$.

Also, the choice of 0 and 1 in the definition of the order property is not particularly important. This is just a convention. What is important is that ϕ separates the values by some positive amount. Say there are $r_1 < r_2$ such that

$$\phi(a_i, a_j) \leq r_1 \text{ when } i < j$$

and

$$\phi(a_i, a_j) \geq r_2 \text{ when } i \geq j.$$

A definition based on this more general situation would be equivalent to the above definition, since we could compose this ϕ with a continuous function $u : \mathbb{R} \to \mathbb{R}$ which sends $(-\infty, r_1]$ to 0 and $[r_2, \infty)$ to 1.

We will check that being unstable is the same as having the order property. The following definition will make it easier to discuss the proof of this claim.

Definition 4.0.4. We say a type p in M is **finitely determined** if for every formula $\phi(x, y)$ and every $\varepsilon > 0$, there is a finite $B \subseteq M$ and some $\delta > 0$ such that for all $c, c' \in M$, if

$$\max_{b \in B} |\phi(b, c) - \phi(b, c')| < \delta$$

then

$$|\phi(p, c) - \phi(p, c')| \leq \varepsilon.$$

Proposition 4.0.5. *T is unstable if and only if it has the order property.*

Proof. ($\leftarrow$) Suppose T has the order property, witnessed by the formula ϕ and sequence $(a_n)_{n \in \mathbb{N}}$ in some M. Let λ be any infinite cardinal. Let μ be the smallest cardinal such that

$$2^{\mu} > \lambda.$$

Let $2^{<\mu}$ denote sequences over $\{0, 1\}$ of length strictly less than μ, ordered lexicographically. Note that $2^{<\mu} \leq \lambda$ by choice of μ.

Recalling that the a_n are ε-separated for some $\varepsilon > 0$, we can take ultrapowers of M to find arbitrarily long sequences witnessing the order property with ϕ. In particular, we can find a bounded sequence $(b_i)_{i \in 2^{<\mu}}$ in some ultrapower $M^{\mathcal{U}}$ such that

$$\phi(b_i, b_j) = \begin{cases} 0 & \text{when } i < j \\ 1 & \text{when } i \geq j \end{cases}.$$

This implies the existence of a type for each cut in

$$B = \{b_i : i \in 2^{<\mu}\}.$$

Any pair of such types has distance at least ε. So, we have a model $M^{\mathcal{U}}$ of T with a set B of cardinality $\leq \lambda$ over which the type space has density $> \lambda$.

($\rightarrow$) Suppose T is unstable. Let λ be a cardinal with

$$\lambda^{\text{density}(L)} = \lambda.$$

If for all $M \models T$, all types over parameter sets of cardinality $\leq \lambda$ in M were finitely determined, then there would be at most $\lambda^{\text{density}(L)} = \lambda$ many types over any such parameter set and T would be λ-stable. So, there is some $M \models T$ with a type p which is not finitely determined, witnessed by a formula $\phi(x, y)$ and some $\varepsilon > 0$. We have that for all finite $B \subseteq M$ and all $\delta > 0$, there are $c, c' \in M$ such that

$$\max_{b \in B} |\phi(b, c) - \phi(b, c')| < \delta$$

and yet

$$|\phi(p, c) - \phi(p, c')| > \varepsilon.$$

We inductively define sequences a_n, b_n, and c_n, as well as sets B_n as follows. Let $B_0 = \emptyset$. Given B_n, choose b_n and c_n so that

$$\max_{b \in B_n} |\phi(b, b_n) - \phi(b, c_n)| < \frac{1}{2}\varepsilon$$

but also

$$|\phi(p, b_n) - \phi(p, c_n)| > \varepsilon.$$

Choose a_n realizing the formulas

$$\phi(x, b_i) \approx \phi(p, b_i)$$

$$\phi(x, c_i) \approx \phi(p, c_i)$$

for all $i \leq n$, which is possible since p is a type in M. Define B_{n+1} by

$$B_{n+1} = B_n \cup \{a_n, b_n, c_n\}.$$

This construction is such that whenever $i < j$ we get

$$|\phi(a_i, b_j) - \phi(a_i, c_j)| < \frac{1}{2}\varepsilon$$

since $a_i \in B_j$ (and hence is considered in $\max_{b \in B_n}$), and whenever $i \geq j$, we get

$$|\phi(a_i, b_j) - \phi(a_i, c_j)| > \varepsilon$$

by choice of a_i. This means that the formula $\psi(x_1, y_1, z_1, x_2, y_2, z_2)$ given by

$$|\phi(x_1, y_2) - \phi(x_1, z_2)|$$

witnesses the order property, using the sequence $(a_n, b_n, c_n)_{n \in \mathbb{N}}$. □

The following examples discuss two simple cases. The application section of this thesis involves proving the instability of some more interesting structures.

Example 4.0.6. If the sorts of M are proper, then $\mathrm{Th}(M)$ is stable. This is because bounded subsets of proper spaces are compact, and it is impossible to find an infinite, ε-separated sequence in a compact space. The theory has sentences which ensure that any model has a finite bound on the length of ε-separated sequences in any given radius ball. In particular, the theory of the real field with its usual metric $(\mathbb{R}, d, 0, 1, +, \cdot)$ is stable in this setting.

Example 4.0.7. Let M be $(\mathbb{N}, d, R)$ where d is the discrete metric and R is a binary relation defined by $R(a, b) = 0$ if $a \leq b$ and $R(a, b) = 1$ otherwise. The relation R is controlled since it is bounded and uniformly continuous. The theory of M is unstable, as witnessed by $R(x, y)$ and the infinite sequence $(0, 1, 2, \dots)$.

Later we will be interested in showing that certain structures arising as ultraproducts are unstable. We will need the following generalization of the order property. The idea is to capture how a sequence of structures may build up to the order property as we go along.

Definition 4.0.8. We say a sequence M_n of L-structures **approaches the order property** if there is an L-formula $\phi(x, y)$ with finite variable tuples x, y of the same length and there is a real tuple r of the same length as x and

y, such that for all $\varepsilon > 0$ and $k \in \mathbb{N}$, for all but possibly finitely many of the M_n, there are $a_1, \ldots, a_k$ in $B_r(M_n)$ such that

$$\phi(a_i, a_j) \leq \varepsilon \text{ when } i < j \leq k$$

and

$$\phi(a_i, a_j) \geq 1 - \varepsilon \text{ when } k \geq i \geq j.$$

Similarly to the earlier order property, this definition would be equivalent if we require instead that $\phi(a_i, a_j) = 0$ when $i < j \leq k$ and $\phi(a_i, a_j) = 1$ when $k \geq i \geq j$, because we can compose with a connective to accomplish this.

An easy example of approaching the order property is the sequence of spaces $M_n = \{1, \ldots, n\}$ with the discrete metric and a relation R with $R(a, b) = 0$ when $a < b$ in the usual sense and $R(a, b) = 1$ otherwise. An interesting observation about this sequence of spaces is that any nonprincipal ultraproduct will be unstable, which can be seen by using R to demonstrate the order property in the ultraproduct. Similarly, if we find an ultraproduct which has the order property, we can infer the existence of a subsequence of its factors which approaches the order property.

The following theorem appears as Theorem 6.1 of [11], but we have changed the wording here.

Proposition 4.0.9. *Assume the continuum hypothesis fails. Let $(M_n)_{n \in \mathbb{N}}$ be a sequence of L-structures which approaches the order property and such that each model has cardinality $\leq 2^{\aleph_0}$. Then there are $2^{2^{\aleph_0}}$ many mutually non-isometric structures of the form $\prod_{\mathcal{U}} M_n$ where $\mathcal{U}$ ranges over nonprincipal ultrafilters over $\mathbb{N}$.*

Proof. We will briefly summarize the argument here, but refer to [11] for the full proof. Say ϕ orders a sequence when that sequence behaves as in the definition of the order property or approaching the order property. This terminology is just to account for the fact that our ϕ is real-valued rather than a true-false relation in the usual sense of an ordering.

Let ϕ be the formula witnessing that $(M_n)_{n \in \mathbb{N}}$ approaches the order property. Without loss of generality, we can assume that for each n, M_n contains a sequence $a_1, \ldots, a_n$ as in the definition of approaching the order property. Note that any ultraproduct over the M_n which uses a nonprincipal ultrafilter over $\mathbb{N}$ will have density $2^{\aleph_0}$, due to the ε-separation of witnessing sequences in M_n.

The proof is a pigeonhole argument that there must be many isometry types among the ultraproducts. This is because there are $2^{2^{\aleph_0}}$ invariants associated to linear orders which will be represented within the ultraproducts using ϕ as the ultrafilter is varied, but for any given ultraproduct, its density bounds how many can be associated to that ultraproduct. The proof requires establishing the existence of these linear orders, showing that by varying the ultrafilter they can be found among the ultraproducts, and proving the required bound. $\square$

Part 2: Application to Metric Geometry

5. Geodesics and CAT(κ) Spaces

In this section, we start applying continuous model theory to metric geometry. The spaces we are interested in have a notion of bounded curvature defined in terms of the behavior of geodesic triangles. We will show that CAT(κ) spaces form an elementary class, and we will show that geodesics are definable in this class.

A standard reference for the geometry material here is Bridson and Haefliger's book [8].

5.1. **Geodesics.** Lines are fundamental in euclidean geometry. Geodesics are one way of extended the notion of lines to more general metric spaces, and we will see that they continue to have an important role. Here we are using a definition which emphasizes a more global, metric approach, rather than appealing to some local, differential structure.

Definition 5.1.1. Let (X, d) be a metric space. A **geodesic** in X is an isometric embedding $\gamma : I \to X$ where I is an interval of $\mathbb{R}$. We call attention to three kinds of geodesics.

(1) If I is a closed interval $[0, \ell]$ we call γ a **(geodesic) segment** from $\gamma(0)$ to $\gamma(\ell)$. We call $\gamma(0)$ and $\gamma(\ell)$ the **endpoints** of γ and say γ has **length** ℓ.
(2) If I is of the form $[0, \infty)$ we call γ a **(geodesic) ray** starting at $\gamma(0)$.
(3) If I is $\mathbb{R}$ we call γ a **(geodesic) line**.

In this document we will usually just say segment, ray, and line.

If a geodesic has $p = \gamma(r)$ for some r, we say p **is on** γ, or that γ **passes through** p, or write $p \in \gamma$.

As was suggested, in a space like $\mathbb{R}^2$, the geodesics correspond to the lines we are used to. A simple case illustrating the more general notion is the 2-sphere with a metric given by angular difference measured from the center. There, the geodesic segments correspond to the segments of a great arcs between two points on the surface. There are no rays or lines in this example.

Definition 5.1.2. A **geodesic triangle** in (X, d) is a triple of geodesic segments $\gamma_1, \gamma_2, \gamma_3$ with endpoints a, b, and b, c, and c, a respectively.

We call $\gamma_1, \gamma_2, \gamma_3$ the **edges** or **sides** of the triangle, and call a, b, c the **vertices** of the triangle.

If $p \in \gamma$ for one of the edges γ of a triangle Δ, we will say p **is on** $\triangle$, or write $p \in \triangle$.

We will see that a metric-based approach to curvature can be carried out using geodesic segments. This is done by examining the behavior of closed loops of geodesic segments. Continuing the sphere example, the geodesic triangles traced on the surface of a 2-sphere have different behavior than those found in the euclidean plane. For example, the angles, areas, and lengths between points on different sides all behave differently.

Not all metric spaces have nontrivial geodesics. There are obvious degenerate cases like a single point or discrete metric spaces which clearly contain no embedded real intervals. But connected spaces can also fail to have geodesics. For example, suppose we give $\mathbb{R}$ the metric

$$d(x, y) = \sqrt{|x - y|}$$

and consider the points 0 and 1. The distance between them is 1, so a geodesic segment between them would be an isometry γ from the usual euclidean $[0, 1]$ to this space. But then $z = \gamma(\frac{1}{2})$ would be a point with

$$d(0, z) = d(z, 1) = \frac{1}{2}.$$

This is impossible with this d, since it means

$$\sqrt{|z|} = \sqrt{|1 - z|} = \frac{1}{2},$$

i.e. both $|z| = \frac{1}{4}$ and $|1 - z| = \frac{1}{4}$.

It is also easy to construct examples where geodesics exist, but only between points that are sufficiently close. For example, this happens in bounded spaces, but also in spaces like euclidean $\mathbb{R}$ with the integers removed.

Definition 5.1.3. Let $D \in \mathbb{R} \cup \{\infty\}$. We say (X, d) is a D-**geodesic space** if for every $x_1, x_2 \in X$ with $d(x_1, x_2) < D$, there is a geodesic segment with endpoints x_1, x_2.

We say (X, d) is a **geodesic space** if every $x_1, x_2 \in X$ are the endpoints of some geodesic segment.

Also note that in general, there could be multiple geodesics between a given pair of points. That is, the endpoints might not determine the segment. For example, opposite poles of a 2-sphere have infinitely many geodesics between them.

5.2. **Model spaces.** We will quantify curvature of geodesic spaces by comparing the behavior of geodesic triangles to those in some special, fixed spaces. These are the classical spaces of constant curvature from differential geometry: hyperbolic spaces, euclidean spaces, and spheres. We will only discuss the 2-dimensional spaces, since we will only need those for our purposes.

Definition 5.2.1. We define the **model space** $\mathcal{M}^2_\kappa$ for $\kappa \in \mathbb{R}$ as follows.

The space $\mathcal{M}^2_0$ is $\mathbb{R}^2$ with the usual euclidean metric.

The space $\mathcal{M}^2_1$ is the unit sphere in $\mathbb{R}^3$ with metric given by

$$d(x, y) = \cos^{-1}(x \cdot y)$$

where $\cdot$ is the usual scalar product on $\mathbb{R}^3$. The images of geodesics in this space are given by intersecting (euclidean) planes through the origin of $\mathbb{R}^3$

with the unit sphere, and the angle between such planes is taken to be the angle between corresponding geodesic segments.

For any $\kappa > 0$, the space $\mathcal{M}^2_\kappa$ is the space obtained from $\mathcal{M}^2_1$ by multiplying distances by $1/\sqrt{\kappa}$.

The space $\mathcal{M}^2_{-1}$ is the upper sheet of the hyperboloid

$$\{x \in \mathbb{R}^3 : \langle x, x\rangle = -1\}$$

with metric given by

$$d(x, y) = \cosh^{-1}(\langle x, y\rangle)$$

where $\langle , \rangle$ denotes the scalar product given by

$$\langle x, y\rangle = x_1y_1 + x_2y_2 - x_3y_3.$$

Again, the geodesics in this space are given by intersecting (euclidean) planes through the origin of $\mathbb{R}^3$ with the sheet, and this can be used to define an angle between geodesic segments.

For any $\kappa < 0$, the space $\mathcal{M}^2_\kappa$ is the space obtained from $\mathcal{M}^2_{-1}$ by multiplying distances by $1/\sqrt{-\kappa}$.

We summarize some properties of these spaces, which can be found throughout section I.2 of [8].

Proposition 5.2.2. *Let $\kappa \in \mathbb{R}$. Let $\mathcal{M} = \mathcal{M}^2_\kappa$.*

(1) The diameter of $\mathcal{M}$ is denoted D_κ and is equal to $\pi/\sqrt{\kappa}$ for $\kappa > 0$ and ∞ otherwise.
(2) $\mathcal{M}$ is a geodesic space. If $x, y \in \mathcal{M}$ have $d(x, y) \le D_\kappa$, then there is a unique geodesic with endpoints x, y, and if B is a closed ball in $\mathcal{M}$ with radius $\le \frac{1}{2}D_\kappa$, then B is convex.
(3) $\mathcal{M}$ has a law of cosines. Given a geodesic triangle with lengths a, b, c and angle θ opposite the side with length c, we have the following.
   (a) When $k = 0$,
   $$c^2 = a^2 + b^2 - 2ab\cos(\theta).$$
   (b) When $k < 0$,
   $$\cosh(\sqrt{-\kappa}c) = C - S$$
   where
   $$\begin{aligned} C &= \cosh(\sqrt{-\kappa}a)\cosh(\sqrt{-\kappa}b) \\ S &= \sinh(\sqrt{-\kappa}a)\sinh(\sqrt{-\kappa}b)\cos(\theta). \end{aligned}$$
   (c) When $k > 0$,
   $$\cos(\sqrt{\kappa}c) = C + S$$
   where
   $$\begin{aligned} C &= \cos(\sqrt{\kappa}a)\cos(\sqrt{\kappa}b) \\ S &= \sin(\sqrt{\kappa}a)\sin(\sqrt{\kappa}b)\cos(\theta). \end{aligned}$$

(4) $\mathcal{M}$ embeds certain triples. Let x_1, x_2, x_3 be three points in an arbitrary metric space (X, d_X) with

$$\sum_{i<j} d_X(x_i, x_j) < 2D_\kappa.$$

Then there are three points $\bar{x}_1, \bar{x}_2, \bar{x}_3$ in $\mathcal{M}$ such that

$$d(\bar{x}_i, \bar{x}_j) = d_X(x_i, x_j)$$

for $i < j$.

These properties are what let us compare geodesic triangles in other spaces to those in the model space.

Definition 5.2.3. A geodesic triangle formed by the corresponding points $\bar{x}_i$ in (4) above, when it exists, is called a **comparison** triangle for the points x_i.

One might wonder why we have defined this as a comparison triangle rather than a comparison triple. Up to isometry, the geodesic triangles in model spaces are determined by the distances between their vertices.

Proposition 5.2.4. *Comparison triangles in $\mathcal{M}^2_\kappa$ for the same triple x_i from (X, d) are unique up to an isometry of $\mathcal{M}^2_\kappa$.*

Proof. This is the second part of Lemma I.2.14 in [8]. □

The map $x_i \mapsto \bar{x}_i$ between the vertices of a triangle and its comparison triangle induces a bijection between points in the segments. This lets us define comparison points for the entire triangle, not just the vertices. Roughly, this is done by matching up points which are the same distance along the corresponding segments.

Definition 5.2.5. Let $\triangle$ be a geodesic triangle in (X, d) with vertices a, b, c. Let $\bar{\triangle}$ be the comparison triangle in $\mathcal{M}^2_\kappa$. Let γ be the side of $\triangle$ from a to b. Let j be the isometry taking γ to the side $\bar{\gamma}$ of $\bar{\triangle}$ from $\bar{a}$ to $\bar{b}$. For $p \in \gamma$, we call $\bar{p} = j(p)$ the **comparison point** for p.

Critically, while the correspondence between points and comparison points is a bijection, and even an isometry when restricted to each side independently, it is not generally an isometry between the geodesic triangle and the comparison triangle. This is because points on different edges may fail to map isometrically. This deviation from an isometry is the fundamental measurement we make when looking at comparison triangles. The way that this deviation varies across the $\mathcal{M}^2_\kappa$ is what lets us quantify the curvature.

5.3. **CAT(κ) spaces.** These are the spaces with a notion of bounded curvature, understood via comparison triangles. Definitions and basic properties of these spaces can be found in the first few sections of part II of [8].

Definition 5.3.1. Let $\triangle$ be a geodesic triangle in a metric space (X, d_X). Let $\bar{\triangle}_\kappa$ be its comparison triangle in $\mathcal{M}^2_\kappa$. We say $\triangle$ satisfies the **CAT(κ) inequality** if whenever $p, q \in \triangle$, we have $d_X(p,q) \leq d_{\mathcal{M}^2_\kappa}(\bar{p}, \bar{q})$ where $\bar{p}, \bar{q}$ are the comparison points in $\bar{\triangle}_\kappa$.

For example, triangles on a standard sphere do not satisfy the CAT(0) inequality, because the distances from one side to another are larger than corresponding lengths in $\mathbb{R}^2$. On the other hand, triangles in the hyperbolic plane do satisfy the CAT(0) inequality, because the relevant distances are shorter than in $\mathbb{R}^2$. This is a formalization of the idea that triangles in spheres are "fatter" than euclidean triangles, and hypebolic triangles are "thinner" than euclidean triangles.

Definition 5.3.2. (X, d) is a **CAT(κ) space** if it is a D_κ-geodesic space where every geodesic triangle whose segment lengths sum to $< D_\kappa$ satisfies the CAT(κ) inequality

Immediate examples are given by the model spaces, and subspaces of them which are closed under geodesics. All triangles in $\mathcal{M}^2_\kappa$ trivially satisfy the CAT(κ) inequality, and moreover can be seen to satisfy the CAT(κ') inequality for $\kappa' \geq \kappa$ by considering the law of cosines in the different model spaces. In more interesting cases, (X, d) will have triangles which vary in what CAT(κ) inequalities they satisfy. As such, knowing that (X, d) is a CAT(κ) space is valuable as a bound rather than complete description of the curvature. It tells us that the triangles in X are "thin enough" for certain arguments.

The existence of comparison triangles lets us define a notion of angle.

Definition 5.3.3. Given three points x, y, z in a CAT(κ) space, we define $\tilde{\angle}_x(y, z)$ to be the (unsigned) angle in the comparison triangle (in $\mathcal{M}^2_\kappa$) at $\bar{x}$. We call this the **comparison angle**.

Note that the comparison angle is not a local property of the sides meeting at x. That is, it depends on the lengths of the segments, and considering other triangles which share initial segments at x may lead to different comparison angles.

This leads to another notion of angle. In any particular CAT(κ) space, given y' and z' on the segments from x to y and from x to z respectively, we have

$$\tilde{\angle}_x(y', z') \leq \tilde{\angle}_x(y, z)$$

as a consequence of the cosine law. This means the limit as y' and z' both approach x is defined.

Definition 5.3.4. Given three points x, y, z in a CAT(κ) space with

$$d(x, y) + d(y, z) + d(z, x) < 2D_\kappa$$

we define $\angle_x(y, z)$ to be $\lim_{y', z' \to x} \tilde{\angle}_x(y', z')$. We call this the **initial angle**.

Furthermore, we clearly always have the following.

Proposition 5.3.5. *Given three points x, y, z in a CAT(κ) space, we have $\angle_x(y,z) \leq \tilde{\angle}_x(y,z)$.*

We will need the following result later when we work with euclidean buildings.

Proposition 5.3.6. ***(Convexity in CAT(0) spaces)*** *Let X be CAT(0). For any pair of segments $c : [0,\ell] \to X$ and $c' : [0,\ell'] \to X$, and for all $t \in [0,1]$, we have*

$$d\left(c(t\ell), c'(t\ell')\right) \leq (1-t)\, d\left(c(0), c'(0)\right) + t d\left(c(\ell), c'(\ell')\right).$$

Proof. This is Proposition II.2.2 of [8]. The proof is a straightforward consequence of euclidean geometry and the CAT(0) inequality. □

5.4. **Approximate midpoints.** In this section we will see how being a geodesic space is equivalent to always being able to find points arbitrarily close to half-way between pairs of points, at least for complete metric spaces. This lets us see another characterization of CAT(κ) spaces which is more amenable to axiomatization in our logic.

Definition 5.4.1. We say m is a **midpoint** of $a, b \in X$ if

$$d(a,m) = d(m,b) = \frac{1}{2} d(a,b).$$

For any $\varepsilon > 0$, we say m_ε is an **ε-approximate midpoint** of $a, b \in X$ if

$$\max(d(a,m_\varepsilon), d(m_\varepsilon, b)) \leq \frac{1}{2} d(a,b) + \varepsilon.$$

Proposition 5.4.2. *Let (X,d) be a complete metric space. The following are equivalent.*

(1) Every pair of points in X has a midpoint.
(2) X is a geodesic space.

Proof. For $(2 \to 1)$, take the point halfway along the geodesic. Now we check $(1 \to 2)$.

Let $a, b \in X$. Let $\ell = d(a,b)$. Define f on $[0,\ell]$ inductively as follows. Let $f(0) = a$ and $f(\ell) = b$. Suppose f is defined at $q_0, q_1, \ldots, q_n$ and already isometric on its defined points. Then define $f(q_{i+1/2})$ to be the midpoint of $f(q_i), f(q_{i+1})$ for each $i = 0, \ldots, n-1$.

We check that f has remained an isometry. It clearly is isometric on

$$\{q_i, q_{i+\frac{1}{2}}, q_{i+1}\}.$$

For each $i = 0, \ldots, n-1$ and $j = 0, \ldots, n$ we can see that

$$d(f(q_j), f(q_{i+1/2})) \leq \left| q_j - q_{i+\frac{1}{2}} \right|$$

by considering the triangle inequality with third point $f(q_i)$ if $q_j \leq q_i$, or third point $f(q_{i+1})$ if $q_j \geq q_{i+1}$. We want to show that this is actually equality. If we were to have strict inequality

$$d(f(q_j), f(q_{i+1/2})) < \left|q_j - q_{i+\frac{1}{2}}\right|$$

we would reach a contradiction by showing $d(a, b) < \ell$ using repeated applications of the triangle inequality on images of

$$0, q_j, q_{i+\frac{1}{2}}, q_{i+1}, \ell$$

if $q_j < q_{i+1/2}$, or images of

$$0, q_i, q_{i+\frac{1}{2}}, q_j, \ell$$

if $q_j > q_{i+1/2}$.

The completeness of X lets us extend f to a function on $I = [0, d(a, b)]$ which remains an isometry, yielding our geodesic. □

Proposition 5.4.3. *Let (X, d) be a geodesic space and $x, y, z \in X$. Then*

$$d(x, y) + d(y, z) = d(x, z)$$

if and only if y is on a geodesic with endpoints x, z.

Proof. If y is on a geodesic γ from x to z, then the equation holds due to γ being an isometry. For the other direction, observe that the construction in the proof above can start with a function f defined on the points

$$0, d(x, y), d(x, z)$$

and iterate from there. □

The following notion is something like a quadrilateral version of comparison triangles. These objects may not always exist, but when they do, they help to simultaneously capture information both about comparison triangles and about sequences of approximate midpoints.

Definition 5.4.4. Let x_1, y_1, x_2, y_2 be four points in a metric space (X, d_X). A **subembedding** in $\mathcal{M}^2_\kappa$ of these points is a four-tuple

$$(\bar{x}_1, \bar{y}_1, \bar{x}_2, \bar{y}_2) \in \mathcal{M}^2_\kappa$$

such that

$$d_X(x_i, y_j) = d(\bar{x}_i, \bar{y}_j) \text{ for } i, j \in \{1, 2\},$$
$$d_X(x_1, x_2) \leq d(\bar{x}_1, \bar{x}_2),$$
$$d_X(y_1, y_2) \leq d(\bar{y}_1, \bar{y}_2).$$

Subembeddings and approximate midpoints can be used to give an alternative characterization of CAT(κ) spaces. This is useful for us, because these notions are easier to work with in continuous logic.

Proposition 5.4.5. *Let (X, d) be a complete metric space. Then the following are equivalent.*

(1) *X is a CAT(κ) space.*
(2) *Every 4 points x_1, x_2, y_1, y_2 in X with*

$$d(x_1, y_1) + d(y_1, x_2) + d(x_2, y_2) + d(y_2, x_1) < 2D_\kappa$$

has a subembedding in $\mathcal{M}^2_\kappa$, and for every $\varepsilon > 0$ and $x, y \in X$ with $d(x, y) < D_\kappa$, there is an ε-approximate midpoint of x, y.

Proof. For the full proof, see Proposition II.1.11 of [8]. We will only comment on the rough idea.

For $(1 \to 2)$, we construct a subembedding by patching together comparison triangles.

For $(2 \to 1)$, we use the subembeddings to show that a sequence of approximate midpoints is cauchy, hence X has midpoints and is a geodesic space. Then the existence of subembeddings can be used again to check the CAT(κ) inequality on geodesic triangles. □

6. Axioms and Quantifiers for CAT(κ) Spaces

6.1. **CAT(κ) spaces form an elementary class.** We can check the elementarity of the class of CAT(κ) spaces for a given κ. The language we will use throughout is just a single sorted language with no additional function or relation symbols.

In the next section, we will produce a set of axioms. For now, we provide a proof using the characterization of elementary classes in Proposition 3.8.2. Interestingly, people working in large-scale geometry seem to be acquainted with both CAT(κ) spaces and ultralimit constructions (what they would call ultraproducts of metric spaces), and make use of the fact that this class is closed under ultraproducts.

Proposition 6.1.1. *Let $\kappa \in \mathbb{R}$. The class of CAT(κ) spaces is elementary.*

Proof. The class is clearly closed under isomorphisms, since these are isometries and must preserve the metric properties involved in the definition of CAT(κ) spaces.

Suppose M is an elementary substructure of a CAT(κ) space N. We will check that M is CAT(κ) by checking that it has approximate midpoints and satisfies the 4-point condition. Existence of approximate midpoints is easily seen by noting that for all $r < r'$, the sentence

$$\sup_{x,y}]_r^{r'} \inf_z]_{r'}^{r'+1} \left((d(x,y) \gtrsim D_\kappa) \cdot \left(\max(d(x,z), d(z,y)) \lesssim \frac{1}{2} d(x,y) \right) \right)$$

is in the theory of N, hence in the theory of M. The $d(x,y) \gtrsim D_\kappa$ term should be dropped if $D_\kappa = \infty$. These imply that all x, y with $d(x,y) < D_\kappa$ in the space have arbitrarily good approximate midpoints. The existence of subembeddings follows trivially from the containment of M in the CAT(κ) space.

Let M be an ultraproduct over $\mathrm{CAT}(\kappa)$ structures M_i. Suppose $a, b \in M$ are such that $d(a,b) < D_\kappa$ and both $d(a,\star)$ and $d(b,\star)$ are $\leq R$. Let $(a_i)_{i\in I}$ and $(b_i)_{i\in I}$ be representatives for a, b respectively. For some large set F in the ultrafilter, each M_i for $i \in F$ satisfies $\phi(a,b)$ with ϕ as in the previous paragraph, since $d(a,b) < D_\kappa$ means $d(a_i,b_i) < D_\kappa$ for these i. The fundamental theorem of ultraproducts gives us $M \models \phi(a,b)$, and so M has approximate midpoints. Let a, b, c, d be 4 points in M relevant for checking the subembedding requirement. A similar argument shows that representatives are such that a_i, b_i, c_i, d_i have a subembedding $\bar{a}_i, \bar{b}_i, \bar{c}_i, \bar{d}_i$ for i in some filter-large set. By homogeneity of the model space, we can assume these subembeddings have one point at the origin and are contained in a closed ball whose radius is at most the largest of $d(a,\star)$, $d(b,\star)$, $d(c,\star)$, and $d(d,\star)$. Since the model space is proper, its ultrapower is itself. So we can compute the ultrapower of the model space $\mathcal{M}^2_\kappa$ alongside our ultraproduct and consider the ultralimits $\bar{a}, \bar{b}, \bar{c}, \bar{d}$ still inside the model space for each sequence $\bar{a}_i, \bar{b}_i, \bar{c}_i, \bar{d}_i$ respectively. By the fundamental theorem of ultraproducts, these resulting limit points serve as a subembedding for a, b, c, d. □

6.2. **Axioms for CAT(κ) spaces.** We will use the characterization of $\mathrm{CAT}(\kappa)$ spaces via the subembedding and approximate midpoint condition to get axioms using only the distance predicate. The idea here is simple, though the axioms themselves may look complicated.

We have two axiom schemas, the first asserts that approximate midpoints exist. This is complicated by the fact that we can only quantify over bounded parts of the space at a time, so we need a sequence of axioms over increasingly large parts of the space.

The second schema asserts that subembeddings exist. This is complicated again by needing to quantify over larger and larger parts of the space. But it is further complicated by the fact that we need to somehow capture quantification over $\mathcal{M}^2_\kappa$ to say that subembeddings exist. Since $\mathcal{M}^2_\kappa$ is proper, its closed balls are compact, and this allows us to finitely describe the distances involved in any ε-net over any ball. We leverage this by asserting that for each $\varepsilon > 0$, any 4-tuple in our space can be matched with a 4-tuple in one of these nets such that the corresponding tuples satisfy the 4-point inequalities to within ε. This is done through a product and sum over permutations of nets, roughly corresponding to what one might expect to be a discrete existential-conjunctive statement. With all of these axioms taken together, the compactness of balls in $\mathcal{M}^2_\kappa$ will ensure that a sequence of these "approximate subembeddings" will converge to a true subembedding.

Note that in these axioms, we will reference distances between points in $\mathcal{M}^2_\kappa$, and the expressions such as $\bar{d}(\sigma(\bar{y}_1), \sigma(\bar{y}_2))$ that appear are actual real numbers, not formulas built from distance and function symbols in our language.

To help prevent some distraction, note that in the quantifiers below we must select bounds, but the particular choice of bounds is not critical. For example, in the $\sup_y]_{n+1}^{n+2}$ appearing in the first axiom schema, the use of $n+1$ and $n+2$ is not special. The bounds are chosen so that we are speaking about a large enough part of the space, possibly relative to earlier quantification. The upper bound is just chosen to be larger than the lower bound for the statement to be saying anything.

We will now define some theories. Immediately after the definition, we show that these theories axiomatize the CAT(κ) spaces. The theories are divided into those with $\kappa \leq 0$ and those with $\kappa > 0$. This is because in the $\kappa > 0$ case, CAT(κ) spaces are only D_κ-geodesic.

Definition 6.2.1. Let L be a language with a single sort. We define L-theories T_κ as follows. First, for $\kappa \leq 0$, T_κ is the theory consisting of the following.

(1) For each $n \in \mathbb{Z}^+$, the sentence

$$\sup_{x,z}]_n^{n+1} \inf_y]_{n+1}^{n+2} \left(\max\left(d(x,y), d(y,z)\right) \lesssim \frac{1}{2} d(x,z) \right).$$

(2) For each $m, n \in \mathbb{Z}^+$ with $m > n$, let $\mathfrak{N}_m$ be a $\frac{1}{2m}$-net of the closed ball of radius n in $\mathcal{M}_\kappa^2$, let $\bar{x}_{1,m}, \bar{y}_{1,m}, \bar{x}_{2,m}, \bar{y}_{2,m}$ be four distinct points in $\mathfrak{N}_m$, let $\bar{d}$ denote the metric in $\mathcal{M}_\kappa^2$, and include the sentence

$$\sup_{x_1,y_1,x_2,y_2}]_n^{n+1} \prod_{\sigma \in S(\mathfrak{N}_m)} \Big(\sum_{i,j \in \{1,2\}} \phi_{\sigma,i,j} + \psi_{\sigma,x} + \psi_{\sigma,y} \Big)$$

where $S(\mathfrak{N}_m)$ is the permutation group of $\mathfrak{N}_m$, the subformula $\phi_{\sigma,i,j}$ is given by

$$d(x_i, y_j) \approx_{\frac{1}{m}} \bar{d}(\sigma(\bar{x}_i), \sigma(\bar{y}_j))$$

the subformula $\psi_{\sigma,x}$ is given by

$$d(x_1, x_2) \lesssim \left(\bar{d}(\sigma(\bar{x}_1), \sigma(\bar{x}_2)) + \frac{1}{m} \right)$$

and the subformula $\psi_{\sigma,y}$ is given by

$$d(y_1, y_2) \lesssim \left(\bar{d}(\sigma(\bar{y}_1), \sigma(\bar{y}_2)) + \frac{1}{m} \right).$$

If $\kappa > 0$, we need to alter the axioms to account for the role of D_κ. So T_κ consists of the following.

(1) In place of (1) above, we use the sentences

$$\sup_{x,z}]_n^{n+1} \inf_y]_{n+1}^{n+2} \Big(d(x,z) \gtrsim (D_\kappa - \frac{1}{n}) \Big) \cdot \Big(\max(d(x,y), d(y,z)) \lesssim \frac{1}{2} d(x,z) \Big).$$

(2) We use the sentences

$$\sup_{x_1,y_1,x_2,y_2}]_n^{n+1} \Big(\psi_{\kappa,n} \Big) \cdot \Big(\prod_{\sigma \in S(\mathfrak{N}_m)} \Big(\sum_{i,j \in \{1,2\}} \phi_{\sigma,i,j} + \psi_{\sigma,x} + \psi_{\sigma,y} \Big) \Big)$$

where $\psi_{\kappa,n}$ is given by

$$(d(x_1,y_1)+d(y_1,x_2)+d(x_2,y_2)+d(y_2,x_1)) \gtrsim (2D_\kappa - \frac{1}{n})$$

and the other subformulas are as defined in (2) above.

Theorem 6.2.2. *Let $\kappa \in \mathbb{R}$ and let $(M,d,\star)$ be any space. Then $M \models T_\kappa$ if and only if M is a CAT(κ) space.*

Proof. Satisfying the first axiom schema, i.e. having each axiom interpret as 0 in the structure, is equivalent to every pair with distance less than D_κ having ε-approximate midpoints for all ε. To see this in the $\kappa > 0$ case, notice that $d(x,y) < D_\kappa$ iff there is some N such that $d(x,y) < D_\kappa - \frac{1}{N}$, hence the first term in the product is bounded away from 0 when $n > N$.

Satisfying the second schema is equivalent to every four points with perimeter $< 2D_\kappa$ having a subembedding in $\mathcal{M}^2_\kappa$. We will just indicate how to parse the axiom to see this. Fix an n and consider the sentences

$$\sup_{x_1,y_1,x_2,y_2}]_n^{n+1} \prod_{\sigma \in S(\mathfrak{N}_m)} \Big(\sum_{i,j\in\{1,2\}} \phi_{\sigma,i,j} + \psi_{\sigma,x} + \psi_{\sigma,y} \Big)$$

with $m > n$. Assume the value of each such sentence is 0 when evaluated in (M,d). Then, since each is a statement about the average supremum of nonnegative formulas over the balls of radius B_n through B_{n+1}, we have that the supremum in $B_n(M)$ is exactly 0. For each m, the product is taken over all permutations of a given $\frac{1}{2m}$-net of $B_n(\mathcal{M}^2_\kappa)$. The sentence evaluating to 0 is equivalent to its product subformula evaluating to 0 for at least one permutation σ. So, at least one $\frac{1}{2m}$-net of $B_n(\mathcal{M}^2_\kappa)$ named by a permutation σ corresponds to a subformula evaluating to 0. Considering the interpretations of $\phi_{\sigma,i,j}, \psi_{\sigma,x}, \psi_{\sigma,y}$, this means that σ provides four points

$$\sigma(\bar{x}_1), \sigma(\bar{y}_1), \sigma(\bar{x}_2), \sigma(\bar{y}_2)$$

in $B_n(\mathcal{M}^2_\kappa)$ which approximately satisfy the subembedding condition, except for an error of $\frac{1}{m}$ in each equality and inequality. However, since this holds for all $m > n$, we obtain four sequences of such points

$$\bar{x}_{1,m}, \bar{y}_{1,m}, \bar{x}_{2,m}, \bar{y}_{2,m}.$$

Since $B_n(\mathcal{M}^2_\kappa)$ is compact, we can choose convergent subsequences to obtain limit points $\bar{x}'_1, \bar{y}'_1, \bar{x}'_2, \bar{y}'_2$ which satisfy the subembedding inequalities exactly.

Conversely, if $(M,d,\star)$ has subembeddings for any four points, we can select approximate witnesses in any $\frac{1}{2m}$-net of the relevant ball by moving $\bar{x}_1, \bar{y}_1, \bar{x}_2, \bar{y}_2$ to the closest point in the net. This possibly changes distances, but only by at most $\frac{1}{m}$ either positively or negatively. So, the equalities and inequalities for the representatives in the net have an error of at most $\frac{1}{m}$. $\square$

6.3. **An alternative approach.** The approach we took above is to characterize CAT(κ) spaces as single-sorted structures, and axiomatize the existence of approximate midpoints and the subembedding condition. The constraints in the subembedding condition come from the possible configurations existing in $\mathcal{M}^2_\kappa$.

Another approach would be to view CAT(κ) spaces as a multi-sorted structure, with one sort being the main space of interest, and an additional sort for the model space $\mathcal{M}^2_\kappa$. Axiomatizing that this additional sort is the model space is straightforward, since $\mathcal{M}^2_\kappa$ is proper.

To axiomatize that the main sort is a CAT(κ) space in this setup, a sup] inf] schema could be used in place of the nets-and-permutations trick we used above. Here the sup would be over 4-tuples of the main sort and the inf over 4-tuples of the $\mathcal{M}^2_\kappa$ sort. This would be a trade-off where we introduce extra sorts to provide a more direct way to talk about subembeddings. This is not a bad trade because of the simplistic behavior of proper spaces. For example, the proper sort (i.e. $\mathcal{M}^2_\kappa$) will be preserved under ultraproducts, elementary substructures, elementary maps, etc.

This alternative approach replaces our original sup] schema with a sup] inf] schema, but in some sense this is irrelevant, because we need a sup] inf] schema for the approximate midpoint axiom in both cases. But it is worth calling attention to these different approaches, since they illustrate a general theme. We are able to take the single-sorted approach because the other space we are "referencing" is proper. Roughly, since sup and inf are inherently approximate notions, we only ever talk about things up to $\varepsilon > 0$, so behavior of compact spaces can be finitely enumerated and described in our theory anyway.

This idea that proper spaces are somehow implicitly accessible can be formally understood by checking that attaching them to structures amounts to an equivalence of categories as in section 3.14.

6.4. **Quantification in geodesic spaces.** This section addresses definability of closed balls at the basepoint when we are dealing with the class of geodesic spaces (which is not an elementary class). This allows us to treat supremums and infimums over such closed balls as part of the language without needing to use integrals in our quantifiers. That is, we effectively have $\sup_{x\in\bar{B}_r}$ and $\inf_{x\in\bar{B}_r}$ quantifiers in our language when working with theories whose models are (at least) geodesic spaces.

The basic observation is that in any geodesic space, we can define the distance to the closed ball of a given radius r simply by taking the distance to the basepoint and subtracting r. That is, we use the formula $d(x,\star) \lesssim r$. As we have seen, the ability to define the distance function is equivalent to having exact quantification over the set.

The important thing to note for us is that this same L-formula interprets as the distance function for every geodesic space. That is, the ball is definable in the class of geodesic spaces. This is important because when we take ultraproducts across such structures, in order to apply Łoś's Theorem, we need a common formula.

These results partially generalize to D-geodesic spaces if we restrict our attention to balls whose radius is strictly smaller than D. The arguments have the same essence as the paragraph above, but the distance function is harder to describe explicitly. That is, one still focuses on the formulas $d(x, \star) \lesssim r$, but generally this is only guaranteed to be the distance function when $||x|| < D$ or equivalently when $(d(x, \star) \lesssim r)$ is less than $D - r$, so we would need to argue further. We are primarily interested in CAT(0) spaces in this thesis, so we will focus on geodesic spaces.

Geodesic spaces are a very convenient case where an obvious formula with the right zero set is also the right formula for measuring distances. Note that the formula we use below will have the right zero set in any structure, whether that structure is a geodesic space or not, but the geometry is critical to understanding the values away from 0 and obtaining definability.

Proposition 6.4.1. *Let L be a language with a single sort. Then the closed balls of any given radius around $\star$ are definable in the class of L-structures which are geodesic spaces$(M, d, \star)$.*

These sets are defined by the formulas $\phi_r(x)$ given by $d(x, \star) \lesssim r$. That is, $\phi_r^M(x)$ is the function $d_M(x, \bar{B}_r(M))$ when interpreted in this class.

Proof. Let $a \in M$ and suppose $\phi_r^M(a) = q$. Notice that $q \in \mathbb{R}_{\geq 0}$. We will check that $d(a, \bar{B}_r(M)) = q$ as well.

If $q = 0$, then $d(a, \star) \leq r$ and hence the following.

$$\begin{aligned} d(a, \bar{B}_r(M)) &= 0 \\ &= q \end{aligned}$$

If $q > 0$, then $d(a, \star) = r + q$. Since M is geodesic, there is a geodesic from $\star$ to a, and on this geodesic there is a point a_r satisfying both of the following.

$$d(\star, a_r) = r$$

$$d(a_r, a) = q$$

If a_r were not the closest point to a among those in $\bar{B}_r(M)$, then the triangle inequality would be violated. Thus,

$$d(a, \bar{B}_r(M)) = q.$$

□

Corollary 6.4.2. *Let $\phi(x, y)$ be an L-formula, and let $r > 0$. Then there are L-formulas equivalent to*

$$\sup_{x \in \bar{B}_r(M)} \phi(x, y)$$

and

$$\inf_{x \in \bar{B}_r(M)} \phi(x, y)$$

for all M in the class of geodesic L-structures.

Proof. The balls are definable in the class and bounded. □

7. Definability of Segments, Rays, and Lines

The purpose of the following sections is to show that geodesic segments, rays, and lines are definable in CAT(κ) spaces using just the distance predicate. Moreover, the definability is across the class of CAT(κ) spaces for fixed, bounded κ. As in the previous section, this lets us quantify over these objects, and having the definability at the class level is critical for maintaining semantics when taking ultraproducts.

There are two different interpretations that could reasonably be called "quantifying over" segments, rays, or lines. One might want to quantify over the points making up a given geodesic segment viewed as a subset of the space. This has a straightforward meaning, but requires either a formula for each geodesic, or a way to parametrize the geodesics within the space.

Alternatively, one might want to quantify over the set of all geodesic segments occuring in the space. If we would like to say the set of segments is a definable subset, we need a superset in mind which is built from sorts. Geodesics are most naturally viewed as isometries from real intervals into the space, or as the images of these functions. But we do not want to include the space of such functions as sorts in our structures.

We can accomplish this by identifying the set of geodesics with a certain definable subset of sequences in the space. Essentially, we can note that any function $\mathbb{R} \to M$ can be thought of as an $\mathbb{R}$-indexed sequence, and then establish the definability of those sequences which represent an isometry. Because definability of a set A of sequences means definability of each set of finite projections $\pi_F(A)$ where F is a finite subset of $\mathbb{R}$, this ultimately reduces the problem to studying the distance of an arbitrary F-tuple to the set of F-tuples which can be sampled from a geodesic. The CAT(κ) inequalities will make it possible to understand this distance function.

The following lemma will be key. Roughly, it generalizes the idea that in CAT(κ) spaces, approximate midpoints are close to actual midpoints, with bounds determined from κ. We obtain the existence of bounds as a consequence of the CAT(κ) inequalities and $\mathcal{M}^2_\kappa$ being proper, but explicit bounds could be obtained using the law of cosines in $\mathcal{M}^2_\kappa$. For example, in CAT(0) spaces, triangles inherit bounds from $\mathbb{R}^2$, where we can appeal to basic euclidean geometry (in particular, Stewart's Theorem or Apollonius's Theorem, or just the cosine law), to compute the required distances between a vertex of a triangle and the opposite side.

The approximate midpoint case is recovered from this lemma when $q = \frac{1}{2}$. The q plays the role of a proportion of the distance along the segment from x to y.

Lemma 7.0.1. ***(Approximate q-point lemma)*** *Let $\kappa \in \mathbb{R}$ and $R > 0$, let $(X, d, \star)$ be a pointed CAT(κ) space.*

Then for all $\varepsilon > 0$, there is $\delta > 0$ such that for all

$$x, y \in \bar{B}_R(X)$$

satisfying

$$d(x, y) < D_\kappa,$$

and for all $q \in [0, 1]$ and all $z \in X$, if z satisfies

$$d(x, z) \leq q \cdot d(x, y) + \delta$$

and

$$d(z, y) \leq (1 - q) \cdot d(x, y) + \delta,$$

then

$$d(z, \gamma_q) \leq \varepsilon$$

where γ_q is the point on the geodesic from x to y at distance $q \cdot d(x, y)$ from x.

Proof. First we show that the claim holds in $\mathcal{M}^2_\kappa$ by contradiction. Suppose there is some $\varepsilon > 0$ where for all $\delta > 0$, there are some $x, y \in X$, $q \in [0, 1]$ and $z \in \mathcal{M}^2_\kappa$ satisfying the hypotheses, but $d(z, \gamma_q) > \varepsilon$. By compactness of $[0, 1]$ and of closed balls in $\mathcal{M}^2_\kappa$, we can select convergent sequences

$$(x_n), (y_n), (q_n), (z_n), (\gamma_{q_n})$$

where for each n these points satisfy the the hypotheses with $\delta = \frac{1}{n}$. Let the limits respectively be

$$x', y', q', z', \gamma_{q'}.$$

Now, note that $\gamma_{q'}$ is actually the point at distance $q' \cdot d(x', y')$ from x' on the geodesic from x' to y'. This is because

$$\begin{aligned} d(x', \gamma_{q'}) &= \lim d(x_n, \gamma_{q_n}) \\ &= \lim q_n \cdot d(x_n, y_n) \\ &= q' \cdot d(x', y') \end{aligned}$$

and similarly

$$d(\gamma_{q'}, y') = (1 - q') \cdot d(x', y').$$

We can also see that

$$d(x', z') \leq q' \cdot d(x', y'),$$
$$d(z', y') \leq (1 - q') \cdot d(x', y'),$$
$$d(z', \gamma_{q'}) \geq \varepsilon.$$

The first two of these inequalities together with the triangle inequality applied to x', y' with third point z' imply that they are actually equalities. So, by uniqueness of the segment from x' to y', we must have $z' = \gamma_{q'}$, contradicting $d(z', \gamma_{q'}) \geq \varepsilon$. This establishes the claim for $\mathcal{M}^2_\kappa$.

Let $(X, d, \star)$ be CAT(κ), and let $\varepsilon > 0$. Let $x, y, q, z \in X$ be as in the hypotheses, and let δ be given by applying the lemma for $\mathcal{M}^2_\kappa$. Let Δ be the comparison triangle in $\mathcal{M}^2_\kappa$ for x, y, z, and recall that this means we have

$$d(\bar{x}, \bar{y}) = d(x, y),$$
$$d(\bar{y}, \bar{z}) = d(y, z),$$
$$d(\bar{x}, \bar{z}) = d(x, z).$$

Consequently, the above argument tells us that

$$d(\bar{z}, \bar{\gamma}_q) \leq \varepsilon$$

where $\bar{\gamma}_q$ is the point $q \cdot d(\bar{x}, \bar{y})$ along the side of Δ from $\bar{x}$ to $\bar{y}$. By the CAT(κ) inequality, this means

$$d(z, \gamma_q) \leq \varepsilon.$$

□

In the next few sections, we will apply this lemma to get definability results.

7.1. **Segments.** The first theorem gives us formulas for the the distance of a point to a given geodesic segment, determined by the endpoints of the geodesic as parameters. In other words, this is definability of the segment from a to b given a, b as parameters.

We will take advantage of the fact that segments are uniquely determined by their (sufficiently close) endpoints in CAT(κ) spaces. Of course, for $\kappa > 0$, we have to account for the role of D_κ. Furthermore, for general continuous logic reasons the formulas depend on how much of the space we expect to talk about at once.

Theorem 7.1.1. ***(Definability of a segment)*** *Let L be any language. For each $\kappa \in \mathbb{R}$ and $R > 0$, there is an L-formula*

$$\phi_{\kappa,R}(x, y, z)$$

such that for any CAT(κ) L-structure M and for all $a, b \in \bar{B}_R(M)$ with $d(a, b) < D_\kappa$, the interpretation

$$\phi^M_{\kappa,R}(a, b, z)$$

is the distance function

$$d(z, \gamma)$$

where γ is the geodesic segment with endpoints a, b.

Proof. Let $R > 0$. Consider the formula $\psi(x, y, z)$ given by

$$d(x, z) + d(z, y) \approx d(x, y).$$

We checked in the section on geodesics that for any such a, b the zero set of $\psi(a, b, z)$ is the segment γ. Unfortunately, ψ does not generally give the distance from z to the segment between a and b (for example, consider triples in $\mathbb{R}^2$), so we instead check that the distance function can be constructed from ψ using Lemma 3.9.6.

Let $\varepsilon > 0$. We will show, independently of M, that there is some $\delta > 0$ such that for all $z \in M$, if $\psi(z) \leq \delta$, then $d(z, \gamma) \leq \varepsilon$. To do this, we will apply our q-point lemma, Lemma 7.0.1.

First, observe that we have the following for any $\delta > 0$. Suppose

$$\psi(a, b, z) \leq \delta$$

with

$$a, b \in \bar{B}_R(M),$$
$$d(a, b) < D_\kappa.$$

By interpreting ψ, we have that

$$d(a, b) - \delta \leq d(a, z) + d(z, b) \leq d(a, b) + \delta.$$

It follows that if $d(a, b) \neq 0$, then

$$\begin{aligned} d(a, z) &\leq d(a, b) - d(z, b) + \delta \\ &\leq (1 - \frac{d(z, b)}{d(a, b)}) \cdot d(a, b) + \delta \\ &\leq (1 - \min(1, \frac{d(z, b)}{d(a, b)})) \cdot d(a, b) + \delta \end{aligned}$$

and

$$d(z, b) \leq \min(1, \frac{d(z, b)}{d(a, b)}) \cdot d(a, b) + \delta.$$

If $d(a, b) = 0$, we can say $\min(1, \frac{d(z,b)}{d(a,b)}) = 1$ and still take these inequalities to hold. In any case, we get both

$$d(a, z) \leq q \cdot d(a, b) + \delta,$$
$$d(z, b) \leq (1 - q) \cdot d(a, b) + \delta$$

where

$$q = \left(1 - \min\left(1, \frac{d(z, b)}{d(a, b)}\right)\right) \in [0, 1].$$

So by 7.0.1, there is some $\delta > 0$ such that $\psi(z) \leq \delta$ implies $d(z, \gamma_q) \leq \varepsilon$, and hence $d(z, \gamma) \leq \varepsilon$.

Now, we can apply the argument in Lemma 3.9.6 independently of the choice of a, b subject to the assumed constraints, since our ε, δ relation does not depend on these choices. This provides a formula $\phi_R(x, y, z)$ which satisfies the conclusion of the theorem. □

Remark 7.1.2. We could obtain explicit bounds in the above proof by considering the geometry of ellipses in $\mathcal{M}^2_\kappa$, since the point $\bar{z}$ must be contained in the ellipse defined by $d(p, \bar{a}) + d(p, \bar{b}) = d(\bar{a}, \bar{b}) + \delta$. For example in $\mathcal{M}^2_\kappa = \mathbb{R}^2$, the semi-major axis of such an ellipse has length $\frac{1}{2}\sqrt{2d(\bar{a}, \bar{b})\delta + \delta^2}$.

The next theorem is another, more direct consequence of Lemma 7.0.1. It says that for any $q \in [0,1]$ and segment, we can definably pick out the point q far along the segment.

Theorem 7.1.3. ***(Definability of the q-point of a segment)*** *Let L be any language. For each $\kappa \in \mathbb{R}$, each $R > 0$ and each $q \in [0,1]$, there is an L-formula*

$$\phi_{\kappa,R,q}(x,y,z)$$

such that for any CAT(κ) L-structure M and any $a, b \in \bar{B}_R(M)$ with $d(a,b) < D_\kappa$, the interpretation

$$\phi^M_{\kappa,R,q}(a,b,z)$$

is the function

$$d(z,\gamma_q)$$

where γ_q is the point at distance $q \cdot d(a,b)$ from a on the geodesic from a to b.

Proof. Consider the formula $\phi(x,y,z)$ given by

$$\Big(d(x,z) \approx q \cdot d(x,y)\Big) + \Big(d(z,y) \approx (1-q) \cdot d(x,y)\Big).$$

When $\phi(x,y,z) \leq \delta$, the hypothesis of the q-point lemma 7.0.1 is satisfied. So the same argument as the last theorem goes through. $\square$

Since there are unique geodesics between points with $d(a,b) < D_\kappa$, we can interpret any pair of such points as endpoints of a geodesic. In other words, the distance of any pair to a pair of end points of a geodesic is given by the 0 function.

This provides a trivial way to quantify over geodesic segments in CAT(κ) spaces by just quantifying over pairs of points. The definability of q-points means that the coordinates of a segment can be referenced, which is what one would expect when quantifying over the set of segments. We will see that things are more interesting in the ray and line case.

7.2. **Rays.** Next we will see how we can work with geodesic rays in CAT(0) spaces if we assume an additional extension property for segments.

Definition 7.2.1. We will say a CAT(0) space has **extensions of segments to rays** if for all segments γ with end points a and b, there is a ray γ' with

$$\gamma'(r) = \gamma(r) \text{ for all } r \in [0, d(a,b)].$$

This is a stronger property than is required in order to get definability of rays. We will comment on weaker sufficient conditions later, but essentially, one can read such conditions off from the requirements for definability. For the spaces we are interested in, the above property will be satisfied and is easier to understand. We will at least provide an example now to suggest why we need some kind of ray-existence assumption.

Example 7.2.2. For each $n \in \{1, 2, \dots\}$ let A_n be a copy of the subspace $[0, n]$ of $\mathbb{R}$, with its usual metric and basepoint $\star = 0$. Consider the sets A_n as being disjoint, and let M' be the disjoint union $\bigcup A_n$. To define a pseudometric on M', it suffices to define $d(p, q)$ where $p \in A_n$ and $q \in A_m$. In this case, let $d(p, q)$ be $d(p, \star) + d(\star, q)$. Finally, let M be the quotient of M' by this pseudometric. The idea is that M is a collection of segments joined at $\star$, and the distances between different segments is defined by the shortest path, which goes through $\star$. Now, notice that there are no rays in M, and in particular, no rays starting at $\star$. However, nonprincipal ultrapowers of M have rays at $\star$. For example, the sequence $(A_1, A_2, \dots)$ yields one. This shows that rays cannot be definable in M.

Handling rays is a more complicated task than dealing with segments. In arbitrary CAT(0) spaces, even with extensions of segments to rays, rays are not necessarily determined by a finite number of points like segments are. Two rays might coincide for some initial length and then branch away from each other. For example, this happens in spaces which are tree-like or derived from path lengths in infinite acyclic graphs.

So, we must view rays as sequences. The most natural choice is as $\mathbb{R}_{\geq 0}$-indexed sequences. We should note that Kleiner and Leeb prove in Lemma 2.4.4 of [14] that in countable ultralimits of CAT(0) spaces, rays (and segments and lines) arise as ultralimits of those objects from the factors. This is nearly the hypothesis of the ultraproduct characterization of definability (prop 3.9.14). But we would need to find a formula whose zero set is the collection of relevant sequences.

Instead, we can adapt the argument for the finite projections, which are easier to handle.

Definition 7.2.3. We will denote the set of rays in M viewed as $\mathbb{R}_{\geq 0}$-indexed sequences by Γ_M.

Proposition 7.2.4. *Let L be any language and let $\kappa \leq 0$. For each finite subset $F \subseteq \mathbb{R}_{\geq 0}$, let $\phi_{\Gamma,F}(x_r : r \in F)$ be the L-formula given by*

$$\sum_{(r,s)\in F^2} \min\left(1, d(x_r, x_s) \approx |r - s|\right).$$

Then in every CAT(0) L-structure with extension of rays to geodesics, the zero set of $\phi^M_{\Gamma,F}$ is the set $\pi_F(\Gamma_M)$.

Proof. Notice that $\phi_{\Gamma,F}(x)$ being 0 is equivalent to $(x_r : r \in F)$ corresponding to a partial isometry

$$F \to M$$

where we view F as a subspace of $\mathbb{R}_{\geq 0}$. Since $F \to M$ is an isometry, it is part of a segment γ. We are assuming segments can be extended to rays, so there is a γ' which restricts to γ and hence to $F \to M$. This verifies the claim. □

We know CAT(0) spaces are an elementary class and in particular this class is closed under ultraproducts. Having extensions of segments to rays is easily seen to be preserved in ultraproducts since segments are definable. So, we can get definability of the set of rays by checking that the zero set of $\phi_{\Gamma,F}$ is preserved under ultraproducts.

Theorem 7.2.5. ***(Definability of the set of real-indexed rays)*** *Let L be any language. The set of $\mathbb{R}_{\geq 0}$-indexed rays is definable in the class of CAT(0) L-structures with extension of segments to rays.*

Proof. Let F be a finite subset of $\mathbb{R}_{\geq 0}$, and let $M = \prod_{\mathcal{U}} M_i$ where each M_i for $i \in I$ is in the class. We need to check that anything in the zero set of $\phi^M_{\Gamma,F}$ arises as an ultralimit of things in $\phi^{M_i}_{\Gamma,F}$. That is, we need to check that any F-projection of a ray in M is the ultralimit of F-projections of rays in the factors M_i.

Write F as $\{r_1, r_2, \ldots, r_N\}$, with $r_i < r_j$ when $i < j$. Let $\gamma \in \pi_F(\Gamma_M)$. That is,

$$\gamma = (\gamma_{r_1}, \gamma_{r_2}, \ldots, \gamma_{r_N})$$

is the F-projection of a ray in M. The points γ_{r_1} and γ_{r_N} are the endpoints of a unique segment in M, and the other points can be viewed as points some proportion of the way along this segment. That is, for each $k \in \{1, \ldots, N\}$ we have the following.

$$d(\gamma_{r_1}, \gamma_{r_k}) = \frac{r_k}{r_N - r_1}$$

$$d(\gamma_{r_k}, \gamma_{r_N}) = 1 - \frac{r_k}{r_N - r_1}$$

For each $k \in \{1, \ldots, N\}$, define q_k by

$$q_k = \frac{r_k}{r_N - r_1}.$$

Since γ_{r_1} and γ_{r_2} are points in the ultraproduct, we can choose representatives (a_i) and (b_i) for their respective classes. In each M_i, the points a_i and b_i determine a segment. We would like to build an element of $\pi_F(\Gamma_{M_i})$ from this segment. However, $d(a_i, b_i)$ might be wrong. That is, we might have $d(a_i, b_i) \neq r_N - r_1$. By possibly extending the segment from a_i to b_i, we can always select a point b'_i on the segment or its extension such that $d(a_i, b'_i) = r_N - r_1$. Notice though that by Łoś's Theorem, for all $\varepsilon > 0$, there is a $\mathcal{U}$-large set G where $d(a_i, b_i)$ is within ε of $d(a, b)$. So we will have $\lim_{\mathcal{U}} d(b_i, b'_i) = 0$.

Now, the segment from a_i to b'_i has points $q_k \cdot d(a, b)$ along this segment for each k. For each $k \in \{1, \ldots, N\}$, let $p_{k,i}$ be that point $q_k \cdot d(a, b)$ far along the segment from a_i to b_i. Notice that because we can extend this segment to a ray, and because we have selected points with appropriate distances,

$$\{p_{k,i} : k \in \{1, \ldots, N\}\}$$

is in $\pi_F(\Gamma_{M_i})$. We will see that the ultralimit of these tuples gives our original segment in the ultraproduct M.

Consider the classes $(p_{k,i} : i \in I)/\sim$ in the ultraproduct M. We already have

$$\begin{aligned}(p_{1,i} : i \in I)/\sim &= (a_i : i \in I)/\sim \\ &= \gamma_{r_1}\end{aligned}$$

by definition of (a_i). We get

$$\begin{aligned}(p_{N,i} : i \in I)/\sim &= (b'_i)/\sim \\ &= (b_i)/\sim \\ &= \gamma_{r_N}\end{aligned}$$

from our earlier note that $\lim_{\mathcal{U}} d(b_i, b'_i) = 0$, and the definition of (b_i). By Łoś's Theorem, the rest of the points are $q_k \cdot d(a, b)$ far along the unique segment from a to b, so we must have

$$(p_{k,i} : i \in I)/\sim = \gamma_{r_k}.$$

We can now conclude the theorem by applying Proposition 3.9.14. □

As a matter of practicality, it can be easier to use $\mathbb{N}$-indexed sequences, because countable index sets seem easier to write controlled functions for. There are a few simple observations that make representing rays with $\mathbb{N}$-indexed sequences possible.

In any CAT(0) space M, there is a map sending each ray

$$\gamma : \mathbb{R}_{\geq 0} \to M$$

to its restriction to $\mathbb{N}$,

$$\gamma \restriction_{\mathbb{N}} : \mathbb{N} \to M$$

which remains an isometry if we view $\mathbb{N}$ as a subspace of $\mathbb{R}$. We can then send this to the sequence

$$(\gamma \restriction_{\mathbb{N}} (n) : n \in \mathbb{N}).$$

The important thing to observe is that the map

$$\gamma \mapsto (\gamma \restriction_{\mathbb{N}} (n) : n \in \mathbb{N})$$

is bijective. This follows from uniqueness of segments given the endpoints in a CAT(0) space.

In short, we can clearly sample the $\mathbb{N}$-indexed sequence from any ray, and if we have an $\mathbb{N}$-indexed sequence which represents an isometry $\mathbb{N} \to M$, then there is a unique way to "fill" this sequence out to a ray $\mathbb{R}_{\geq 0} \to M$. We will record this as a lemma for reference.

Lemma 7.2.6. *In any CAT(0) space M, the map pairing each ray*

$$\gamma : \mathbb{R}_{\geq 0} \to M$$

with its corresponding $\mathbb{N}$*-indexed sequence*

$$(\gamma(n) : n \in \mathbb{N})$$

is a bijection between the set of rays in M *and the set of* $\mathbb{N}$*-indexed sequences representing isometries* $\mathbb{N} \to M$*.*

We can then immediately see the definability of $\mathbb{N}$-indexed rays as a corollary of the $\mathbb{R}_{\geq 0}$-indexed case.

Corollary 7.2.7. (Definability of naturals-indexed rays) *Let* L *be any language. The set of* $\mathbb{N}$*-indexed rays is definable in the class of CAT(0) L-structures with extension of segments to rays.*

Proof. Finite subsets of $\mathbb{N}$ are also finite subsets of $\mathbb{R}_{\geq 0}$. □

We can also explicitly write a formula with the $\mathbb{N}$-indexed rays as the zero set.

Proposition 7.2.8. *Let* L *be any language. Let* $\phi_\Gamma(x_n : n \in \mathbb{N})$ *be the* L*-formula*

$$\sum_{n=0}^{\infty} \left(\sum_{m=0}^{n-1} \max\left(1, d(x_n, x_m) \approx |n-m|\right) \right) \cdot \frac{1}{n} \cdot 2^{-n}.$$

Then in any CAT(0) L-structure M*, the zero set of* ϕ_Γ^M *is always the set of* $\mathbb{N}_{\geq 0}$*-indexed rays in* M*.*

Proof. Notice that $\phi_\Gamma(x) = 0$ iff $d(x_n, x_m) = |n-m|$ for all $n, m \in \mathbb{N}$. This is essentially the definition of an isometry $\mathbb{N} \to M$, so x corresponds to an $\mathbb{N}$-indexed ray. □

Later, we will want to quantify over rays which start at the basepoint of a pointed metric space. This is a special case of the following.

Theorem 7.2.9. ***(Definability of rays starting in definable sets)*** *Let* L *be any language. Let* ϕ *be any* L*-formula such that the zero set of* ϕ *is definable in the class* $\mathcal{C}$ *of CAT(0) L-structures with extension of segments to rays.*

Then the set of $\mathbb{R}_{\geq 0}$*-indexed rays which start in the zero set of* ϕ *is definable in* $\mathcal{C}$*. Similarly for* $\mathbb{N}$*-indexed rays.*

Proof. We will use the ultraproduct characterization of definability again, and see that this is an easy adaptation of the earlier argument. Let $M = \prod_{\mathcal{U}} M_i$. Write Γ_ϕ^M or $\Gamma_\phi^{M_i}$ for the set of rays which start in the zero set of ϕ in M or one of the M_i. Let F be any finite subset of $\mathbb{R}_{\geq 0}$.

If $\gamma \in \pi_F(\Gamma_\phi^M)$, then γ is the F-projection of a ray which starts in the zero set of ϕ^M. Thus, γ extends to an $\{0\} \cup F$-projection

$$\gamma' = (\gamma_0, \gamma_f : f \in F)$$

of a ray where

$$\gamma_0 \in \mathrm{ZeroSet}(\phi^M).$$

Our proof of Theorem 7.2.5 shows that γ' must be the ultralimit of $\{0\} \cup F$-projections of rays in the factors M_i. The only additional thing to note is that because the zero set of ϕ is definable, we can carry out that proof using a representative for γ_0 chosen from among the zero sets ϕ^{M_i}. Note, rather importantly, that we did not need to move the starting points of the segments in that construction. $\square$

Corollary 7.2.10. *We can quantify over rays starting at the basepoint. We will denote this set by $\Gamma_\star$, or by $\Gamma_\star M$ when we have a space M in mind.*

By quantification over the rays or the rays starting in some definable set, we will mean the $\mathbb{N}$-indexed rays and quantification over countable sequences as in 2.10.1.

In particular, if we are only concerned with rays starting at the basepoint $\star \in M$, then each finite projection is bounded, so we can use exact quantifiers. That is, we can just write

$$\inf_{(x_n : n \in \mathbb{N}) \in \Gamma_\star} .$$

We will close this section with comments on the definability of points along a ray. In the last section, we saw the definability of segments (thm 7.1.1) and of q-points (thm 7.1.3) which showed how to quantify over the points in a segment given its end points. If we are given an $\mathbb{N}$-indexed ray $(a_n : n \in \mathbb{N})$, then each pair a_n, a_m is the pair of end points of a segment. So, those results continue to apply for subsegments of rays. In particular, this shows that we do not lose the ability to reference the $\mathbb{R}_{\geq 0}$-coordinates of a ray, even if we only work with $\mathbb{N}$-indexed rays. If we want to talk about $\gamma(r)$, we can pick $n, m \in \mathbb{N}$ such that $r \in [n, m]$, and then use the definability results to talk about $\gamma(r)$ as the point $(r-n)/m$ far along the segment from $\gamma(n)$ to $\gamma(m)$.

7.3. **Lines.** We will start this section with an example showing a limitation with lines that did not happen with rays. The point here is that assuming extensions of segments to lines is not sufficient to get definability of lines through the basepoint.

Example 7.3.1. We will take a subset of $\mathbb{R}^2$ and give it a different metric. Let M be the subset

$$\{(x, 0) : x \in \mathbb{R}\} \cup \{(x, y) : x \in (0, 1], y \in \mathbb{R}\}$$

with basepoint $\star = (0, 0)$, and define the metric on M as follows. Let $p = (x_p, y_p)$ and $q = (x_q, y_q)$. If $x_p = x_q$, then $d(p, q) = |y_p - y_q|$. Otherwise, $d(p, q)$ is

$$|y_p| + |x_p - x_q| + |y_q|.$$

This makes M a CAT(0) space, since all triangles are degenerate. Now, consider lines λ_k corresponding to $\{\frac{1}{k}\} \times \mathbb{R}$, say with $\lambda_k(0) = (\frac{1}{k}, 0)$. For large k, λ_k becomes an arbitrarily good approximation of a line with $\lambda(0) = \star$, and if we take a nonprincipal ultrapower of M, then these lines will in fact yield a line with $\lambda(0) = \star$. However, each λ_k is bounded away from the lines in M which pass through $\star$. So, lines through $\star$ must not be definable in M.

To see where the argument for rays breaks down in this example, notice that if we take a line that is nearly through $\star$, there is no way to guarantee that a nearby segment passes through $\star$. What typically happens is that there are segments from $\star$ to $\lambda(n)$ and from $\star$ to $\lambda(-n)$, but these two segments might not be able to be put together form a line. More elaborate variations of this example could be constructed by moving the basepoint around and taking unions of copies of this example to show lines through any $\bar{B}_r(\star)$ are not definable.

Nevertheless, if we work with CAT(0) spaces with extensions of segments to lines, then we can argue exactly as in the previous section, using the line extensions in place of ray extensions. We will simply record the results for this section.

Theorem 7.3.2. ***(Definability of the set of lines)*** *Let L be any language. The set of $\mathbb{R}$-indexed lines is definable in the class of CAT(0) L-structures with extension of segments to lines. Similarly for $\mathbb{Z}$-indexed lines.*

Also, like our comments at the end of the ray section, the segment definability results let us quantify over subsegments of lines.

7.4. **Comments on the relation between segments, rays, and lines.** It is worth mentioning some alternative approaches to the definability of rays and lines. Above, we viewed rays and lines as sequences of points which are somehow coherent. We could just as well view them as sequences of segments which cohere. This is not substantially different, but calls attention to what rays and lines are from the perspective of the logic.

Another approach to lines would be to think of them as a pair of rays γ_+ and γ_- which share a starting point

$$\gamma_+(0) = \gamma_-(0)$$

and diverge such that

$$d(\gamma_+(r), \gamma_-(r)) = 2r.$$

This view is worth noting because it naturally leads to the question of whether we could define other configurations of rays, for example those sharing a starting point but diverging at some other rate, say

$$d(\gamma_+(r), \gamma_-(r)) = r\sqrt{2}.$$

However, while this approach would work for lines in CAT(0) spaces with extension of segments to lines, we cannot generally get these other configurations of rays. The following example demonstrates the problem.

Example 7.4.1. Construct a metric space M as follows. The idea to have in mind is that M will be an amalgamation of planar sections of various angles, all joined at a single basepoint $\star$. The distances between points on different sheets will "go through" the basepoint. For each $n \in \mathbb{N}_+$, define A_n to be the subspace of $\mathbb{R}^2$ with its usual metric consisting of the points whose polar angles θ satisfy $\theta \in [0, \frac{\pi}{2} - \frac{1}{n}]$ and keep the basepoint $\star = (0,0)$. Now consider the sets A_n as being disjoint. Given $p \in A_n$ and $q \in A_{n+1}$, let their distance be $d(p,\star) + d(\star, q)$. Then let M be the union of these A_n, and quotient by the pseudometric we have just formed in order to identify all of the basepoints at $\star$.

This M is a CAT(0) space, and for every $\varepsilon > 0$ and $N \in \mathbb{N}$, it has pairs of rays γ_1, γ_2 such that $d(\gamma_1(n), \gamma_2(n))$ is within ε of $n\sqrt{2}$ for all $n \leq N$. However, M does not have any pairs of rays such that $d(\gamma_1(n), \gamma_2(n)) = n\sqrt{2}$ for all n. Considering ultrapowers of M shows that such pairs of rays are not definable in any sense in M. As usual, the issue is that pairs of rays can be arbitrarily close to satisfying some formulas, while not needing to be close in the metric to any actual pair satisfying the formulas exactly. More specifically in this example, rays can diverge at approximately the right rate, there are no rays which diverge at the correct rate, let alone nearby rays which do so.

One way to avoid this problem would be to require some sort of "widening" property, or if we knew we were working in spaces where planar sections extended into full planes. We will see later that in some important theories with additional structure like this, pairs of rays with certain growth rates of $d(\gamma_1(n), \gamma_2(n))$ become important features of a space and are definable.

7.5. **Requirements for definability.** In this section we comment on the remark from the ray definability section that extension of segments to rays is a stronger requirement than is needed.

Suppose we are in a CAT(0) space M where $\mathbb{N}$-indexed rays starting at $\star$ are definable. Then we have a formula $\psi(x_n : n \in \mathbb{N})$ for $d'(x, \Gamma_\star)$ in M, where

$$d'(x, \Gamma_\star) = \inf_{\gamma \in \Gamma_\star} \left(\sum \frac{d(x_n, \gamma_n)}{1 + d(x_n, \gamma_n)} \cdot 2^{-n} \right).$$

Note that this term

$$\sum \frac{d(x_n, \gamma_n)}{1 + d(x_n, \gamma_n)} \cdot 2^{-n}$$

defines a metric d' on countable sequences.

Moreover, the definability implies that for all $\varepsilon > 0$, there is $\delta > 0$ such that M satisfies the following sentence

$$\sup_{x \in \prod B_n} \inf_{y \in \prod B_n} (\psi(x) \gtrsim \delta) \cdot ((\psi(y) \approx 0) + (d'(x,y) \lesssim \varepsilon)).$$

Roughly, the meaning of this sentence is that every δ-approximate Γ_*-like sequence x (as measured by ψ) has arbitrarily good approximate Γ_*-like sequences y within distance ε (in the weighted sum metric d' above).

Now, suppose we are in some other CAT(0) space which satisfies this last sentence. Then, for all $\varepsilon > 0$, there is $\delta > 0$ given above such that if x has

$\psi(x) < \delta$ we can construct a cauchy sequence as follows. Pick a sequence ε_n such that

$$D = \sum_{n=0}^{\infty} \varepsilon_n < \infty$$

and inductively select witnesses y_n which have $\psi(y)$ smaller than the corresponding δ_n. This results in a sequence converging to some y with $\psi(y) = 0$ and $d(x, y) \leq \varepsilon + D$. We can then make D as small as we want and carry out this argument to show that there is some y with $\psi(y) = 0$ and $d(x, y) \leq \varepsilon$.

This shows that satisfying the above sentence is enough to be able to get definability using 3.9.6.

There are two things to keep in mind. One is that the zero set of ψ might not interpret as expected in an arbitrary model. Here, if we work over CAT(0) spaces, we know the zero set of ψ corresponds exactly as we intend to the rays starting at $\star$. The second thing is that this axiom encodes the information for a particular ε-δ argument. In general, there may be spaces where the zero set is definable, but according to a different argument using different ε-δ pairs.

7.6. **Comments on the relation between definability and existence.** The existence of certain objects can be axiomatized, but with a caveat. This is interesting because the superficial problem with inf statements is that we only guarantee approximations, not necessarily actual realizations. However, as we have just seen, if we axiomatize ε-δ information, we can use inf statements to ensure cauchy sequences exist and obtain actual realizations of a zero set. The catch is that this ε-δ information is part of the axioms and so must hold for all of the models being axiomatized.

Let us expand on this relation between such axioms and the extension of segments to rays. We will continue with the same notation. Notice that the sentences above do not quite guarantee that segments extend. What they do guarantee is that sufficiently long segments starting at $\star$, which have small ψ values by virtue of being long segments, must have nearby rays. How near the rays must be is controlled by how long the segment is. But, even in a saturated extension, long segments do not necessarily need to extend. Moreover, such long segments might not even exist, and there might not be any rays in the space at all, yet these axioms may be satisfied.

However, we can find axioms which do guarantee that segments extend. For example, we could use axioms which say any segment has arbitrarily good approximate rays extensions, and then encode ε-δ information saying that any two δ-approximations of extensions of a given segment must be within ε of each other. This would force unique extensions to exist.

7.7. **Another way to skin the CAT(κ) axioms.** We will reconsider the CAT(κ) axioms in light of the previous two sections. Recall that we have axiomatized CAT(κ) spaces based on two primary behaviors:

- certain pairs of points have approximate midpoints, and

- certain 4-tuples have subembeddings in $\mathcal{M}^2_\kappa$.

The first of these was unsurprisingly axiomatized using a $\int \sup \int \inf$ schema, because it is just an approximate existence claim. The second required some finesse, because it needs to guarantee the actual existence of something. We ensured this by exploiting the fact that $\mathcal{M}^2_\kappa$ is proper.

However, the previous discussions hint at an alternative view of the axiomatization. The role of the second schema about subembeddings is that it can be used to ensure the following two things.

- For any $\varepsilon > 0$, there is $\delta > 0$ such that δ-good approximate midpoints are ε-close to each other, where the $\varepsilon - \delta$ pairs come from $\mathcal{M}^2_\kappa$. This ensures the existence of unique midpoints and hence the existence of unique geodesic segments.
- By selecting one of the points in the 4-tuple to be on the segment between two others, the subembedding condition degenerates to a $\mathrm{CAT}(\kappa)$ triangle inequality, showing that this inequality is satisfied.

Notice that the ε-δ information for approximate midpoints can be axiomatized as in the discussions from previous sections. This axiomatizes spaces with geodesic segments existing according to the behavior in $\mathcal{M}^2_\kappa$. We could also obtain this by more directly axiomatizing the geodesic segment existence and definability. In any case, this is enough to get definability results like picking out points dyadic distances along segments. Then, in such spaces, the $\mathrm{CAT}(\kappa)$ inequalities for geodesic triangles can be axiomatized by quantifying over triples of points (implicitly corresponding to geodesic triangles), and bounding distances between (dyadic) points on the segments according to $\mathcal{M}^2_\kappa$, say by using the appropriate cosine law.

The end result is the same, but the approach avoids the messy, direct manipulation of finite nets in $\mathcal{M}^2_\kappa$, and suggests a formal characterization of $\mathrm{CAT}(\kappa)$ spaces as those spaces which have geodesics "like in" $\mathcal{M}^2_\kappa$ and triangles "like in" $\mathcal{M}^2_\kappa$. This agrees nicely with the more typical presentation of $\mathrm{CAT}(\kappa)$ spaces in geometry literature as D_κ-geodesic spaces satisfying the $\mathrm{CAT}(\kappa)$ inequalities.

8. Flats and Atlases

So far the discussion has been focused on embeddings of subspaces of $\mathbb{R}$ into a space. In the upcoming sections, we will look at embeddings from $\mathbb{R}^k$ for higher k as well. We will see that the way these embeddings interact with each other can be used to define a class of spaces with richer definable structure, particularly on rays.

8.1. **Flats.**

Definition 8.1.1. A k-**flat** in (X, d) is a subspace isometric to $\mathbb{R}^k$. In particular, 1-flats are the images of geodesic lines. A k-flat is **maximal** if it is not contained in any k'-flat with $k' > k$.

Notice that the definition is of a special kind of subspace, not the isometries themselves. For example, any 2-flat has infinitely many isometries with $\mathbb{R}^2$ witnessing that it is a 2-flat. In particular, we say that the euclidean plane $\mathbb{R}^2$ itself has only one 2-flat. This is in contrast to our definition of geodesics, which kept track of the isometry, though we often speak as though the geodesic is its image for convenience.

More generally, the adjective "flat" is used as a modifier to describe subspaces A of a space X which are the isometric image of some subspace of some $\mathbb{R}^k$. For example, it is common to come across discussion of flat triangles, flat quadrilaterals, flat strips, and so on.

Definition 8.1.2. A flat in (X, d) is **regular** if it is contained in exactly one maximal flat. Otherwise, if the flat is contained in more than one maximal flat, it is called **singular**.

Similarly, we will say that a geodesic $\gamma : I \to X$ is regular when its image is contained in exactly one maximal flat, and otherwise we say it is singular.

Regular and singular flats are in separate orbits under isometries. This is a consequence of the following easy observation.

Proposition 8.1.3. *Let A be a subspace of X, and let $j : X \to Y$ be an isometry. Then for each k, the map j induces a bijection between the k-flats F containing A and the k-flats $j(F)$ containing $j(A)$.*

Even if a pair of k-flats are both regular or both singular, there is not necessarily an isometry carrying one to the other. This means there is some opportunity to further differentiate within regular and singular flats. We will return to this point later, but for now we merely call attention to it.

In later sections, we will look at spaces which are a union of many k-flats, with these k-flats intersecting as smaller k'-flats. The following example demonstrates one of the simplest interesting ways this can happen.

Example 8.1.4. Let X be the subspace of $\mathbb{R}^3$ consisting of the union of the xy-plane P and the z-upper-half of the xz-plane, H. That is, X is

$$\overbrace{\{(x, y, 0) : x, y \in \mathbb{R}\}}^{P} \cup \overbrace{\{(x, 0, z) : x \in \mathbb{R}, z \in (0, \infty)\}}^{H}.$$

Use the usual distance within P and H, but for $(p, h) \in P \times H$, let $d(p, h)$ be

$$\min_{y \in \{(0,y,0): y \in \mathbb{R}\}} d(p, y) + d(y, h)$$

which can be thought of as the path length from p to h when confined to paths in X.

We have defined X as the union of a 2-flat with a half-plane, but we can also view it as the union of two 2-flats. One of these is the plane P. For the other, we can select the union of the y-nonnegative-half of P together with H. That is,

$$P_{y \geq 0} = \{(x, y, 0) : x \in \mathbb{R}, y \geq 0\}$$

together with H.

This pair of 2-flats we have just discussed have the half-plane $P_{y\geq 0}$ as their intersection. The x-axis

$$\{(x, 0, 0 : x \in \mathbb{R}\}$$

is singular, as is any line or ray confined solely to $P_{y\geq 0}$. However, there are regular geodesics as well. For example, the (geodesic) line given by the y-axis

$$\{(0, y, 0) : z \in \mathbb{R}\}$$

exists in only one 2-flat in X, namely the 2-flat P.

There are, of course, other ways to write X as a union of planes or half-planes.

8.2. **Rays in CAT(0) spaces.** Given a CAT(0) space, we can consider the set of rays from the origin and equip it with a metric. One possible metric is determined by the limiting behavior of comparison angles as we lengthen the initial segments of the rays.

Definition 8.2.1. Let $(M, d, \star)$ be a CAT(0) space with basepoint $\star$. We write $\Gamma_\star M$ for the set of geodesic rays starting at $\star$. We define the metric $\angle_\infty$ called the **angle at infinity** on $\Gamma_\star M$ by

$$\angle_\infty(\gamma_1, \gamma_2) = \lim_{n\to\infty} \tilde{\angle}_\star(\gamma_1(n), \gamma_2(n)).$$

This metric is natural and useful in the metric geometry setting, and we will use it to help define some objects as we go along. However, it does not play well with continuous logic, and we will need to address this later.

Example 8.2.2. Let M be the union of three copies of $[0, \infty)$, identified at their 0 points, and equipped with the path length metric. For convenience, call the three copies X, Y, and Z, and refer to their points as $X(t)$ with $t \in [0, \infty)$, etc. Let the basepoint $\star$ for M be $X(1)$.

Consider the rays γ_Y and γ_Z which start at $\star$ and go along the Y branch and Z branch of the space, respectively. That is,

$$\gamma_Y(t) = \begin{cases} \star & \text{when } t = 0 \\ X(1-t) & \text{when } 0 < t < 1 \\ Y(t-1) & \text{when } t \geq 1 \end{cases}$$

and similarly for γ_Z but with $Z(t-1)$ instead. Then

$$\angle_\infty(\gamma_Y, \gamma_Z) = \pi$$

even though γ_Y and γ_Z share an initial segment.

If we continue with this example, we can get a sense of why this metric is problematic in continuous logic. The issue is that pointwise convergence of rays does not imply convergence of the angle at infinity. Knowing about bounded parts of the space does not provide any information.

Example 8.2.3. For each n, define the pointed space M_n obtained by giving M the new basepoint $X(n)$ instead. Informally, we are pushing the branch point 0 farther from the basepoint $\star$. We can find a sequence of pairs of rays $\gamma_{Y,n}$ and $\gamma_{Z,n}$ in M_n such that for each $t \in [0, \infty)$ we have

$$\lim_{n\to\infty} \gamma_{Y,n}(t) = \lim_{n\to\infty} \gamma_{Z,n}(t)$$

but

$$\begin{aligned} \lim_{n\to\infty} \angle_\infty(\gamma_{Y,n}, \gamma_{Z,n}) &= \lim_{n\to\infty} \pi \\ &= \pi. \end{aligned}$$

This shows that we cannot generally have $\angle_\infty$ as a relation on the definable set $\Gamma_\star$ in a CAT(0) metric structure, for example.

The reader might wonder whether focusing on the rays starting from the basepoint neglects much of the structure of the space of all rays (say, with $\angle_\infty$ as a pseudometric). In general spaces, this is the case, but CAT(0) spaces have a nice property which makes $\Gamma_\star$ in a sense representative of the space of all rays. Consider the equivalence relation defined by the following notion.

Definition 8.2.4. Two rays γ_1, γ_2 are **asymptotic** if there exists $K > 0$ such that

$$d(\gamma_1(t), \gamma_2(t)) \leq K$$

for all t.

We denote the set of equivalence classes of rays in M modulo the relation of being asymptotic by $M(\infty)$.

The next proposition shows how rays at $\star$ serve as representatives for the classes in $M(\infty)$.

Proposition 8.2.5. *Let $(M, d, \star)$ be a complete CAT(0) space. For any ray γ, there is a unique ray $\gamma_\star$ starting at $\star$ which is asymptotic to γ.*

Proof. A proof can be found in Proposition II.8.2 of [8]. Briefly, one can show uniqueness by noting that for two rays γ_1, γ_2 starting at $\star$, the function $d(\gamma_1(t), \gamma_2(t))$ is 0 at $t = 0$, bounded, and defined for all $t \geq 0$, but can also shown to be convex in CAT(0) spaces and hence constantly 0. Then one can show that at least one ray at $\star$ asymptotic to γ exists by showing that the sequence of segments γ_n from $\star$ to $\gamma(n)$ converges to a ray. □

The set $M(\infty)$ is usually called the Gromov boundary or Tits boundary, and $\angle_\infty$ is usually defined on it. Often the length metric derived from $\angle_\infty$ is used instead and called the Tits metric.

The intuition is that $M(\infty)$ represents a "boundary at infinity" of M. This view comes from thinking of rays as paths out to infinity determining ideal end points. This construction can be used to form an important compactification of M. We will not discuss that compactification, but this boundary will play an interesting role going forward.

The following proposition becomes relevant later when we discuss spherical buildings.

Proposition 8.2.6. *Let $(M, d, \star)$ be a CAT(0) structure. Then $\Gamma_\star M$ with metric $\angle_\infty$ is a CAT(1) space.*

Proof. See Theorem II.9.13 of [8]. □

A very basic example of this last proposition is the case where M is just a euclidean space.

Proposition 8.2.7. *Let $\mathbb{E}$ be a euclidean space $\mathbb{R}^k$ with $k > 0$. Then $\Gamma_\star \mathbb{E}$ with metric $\angle_\infty$ is isometric to the sphere $\mathbb{S}^{k-1}$ with the angular metric.*

Proof. Notice that the $\angle_\infty$ distance between rays from the origin in a euclidean space is just their initial angle. The map sending $\gamma \in \Gamma_\star \mathbb{E}$ to $\gamma(1)$ gives the required isometry. □

8.3. **Projective Geometry.** In the remaining sections, we will look at objects which generalize projective planes, and will be interested in interpreting these structures in certain CAT(0) spaces. This section is intended to orient intuition.

We will sketch an argument that $(\mathbb{R}^3, d, 0)$ interprets the projective plane and the incidence geometry on it. From this, given some parameters, we can interpret the real field using classical constructions. There are more direct ways to get the real field in this setting than the constructions we are about to use. Consequently, this exercise will seem strange to readers who have absorbed the impact of knowing that both $(\mathbb{R}^3, d)$ and the real field are proper spaces. Nevertheless, the sketch helps provide a frame of reference for the arguments we will make later in more general settings and with different objects. The important thing to take away is that we use definable sequences to obtain an incidence structure and interpret a field.

Let V_1 be the set of lines through the origin of $\mathbb{R}^3$, and let V_2 be the set of planes through the origin. The real projective plane can be constructed as V_1, with the projective points being the elements of V_1, and the projective lines being the sets

$$\{\ell \in V_1 : \ell \subseteq P\}$$

as P varies members of V_2. The incidence relation is defined by saying that a projective point ℓ is on a projective line P iff $\ell \subseteq P$.

It is easy to check that for any two distinct projective points, there is exactly one projective line which they both lie on, and for any two distinct projective lines, there is exactly one projective point which lies on both. This statement corresponds to the fact in $\mathbb{R}^3$ that any two distinct lines through the origin determine a unique plane through the origin, and any two planes through the origin determine a unique line through the origin. This projective plane enjoys some additional properties, e.g. Desargues's Theorem.

Now we will focus on definability in $(\mathbb{R}^3, d, 0)$. We can define the set of lines through the origin as a subset of

$$\{\star\} \times \bar{B}_1^2 \times \bar{B}_2^2 \times \bar{B}_3^2 \times \cdots$$

We can view this set of lines as a copy of the real projective plane.

To get the incidence geometry, we still need a copy of the set of projective lines, i.e. planes through the origin. One option is to generalize the approach we took for lines by viewing planes first as embeddings j of $\mathbb{Z}^2$. We would then select an appropriate indexing and choose a formula to measure the error of countably many variables from being an isometry from the relevant space and having $j(0) = 0$. This approach works easily in $\mathbb{R}^3$ where the fact that the space is proper can be used.

As an alternative approach, we could capture the planes by viewing them as pairs of lines at some pre-specified angle. This requires choosing a formula that measures the error of the distance of two lines through the origin from being lines with such an angle. Again, this approach works easily in $\mathbb{R}^3$ because of properness and the continuity of angles with respect to the weighted sum metric we have on lines.

In any case, we can define a set and view it as a copy of the set of planes in $\mathbb{R}^3$ through the origin, i.e. the projective lines. Also note that in either case, we have access to "axes" for each plane. These are either pre-specified coordinates of the countable sequence, or the coordinates of the pair of lines (which are really countable sequences).

Now, we need the incidence relation. We will see that there is a formula $\phi(x, y)$ with x ranging over projective points (the definable set corresponding to lines in $\mathbb{R}^3$ through the origin) and y over the projective lines (the definable set corresponding to planes in $\mathbb{R}^3$ through the origin), such that $\phi(x, y)$ is 0 exactly when x lies on y. This zero set is definable due to compactness of balls in $\mathbb{R}^3$.

Thus, given a projective point a, we can quantify over those y with $\phi(a, y) = 0$. Similarly, given a projective line b, we can quantify over those x with $\phi(x, b) = 0$. We can quantify over incident pairs, as well as "the projective points on this given line" and "the projective lines which pass through this given point".

The formula $\phi(x, y)$ can be constructed as follows. Let ℓ_1, ℓ_2, ℓ_3 be lines through the origin in $\mathbb{R}^3$ with $\ell_2 \neq \ell_3$, and consider them as maps from $\mathbb{Z} \to \mathbb{R}^3$ with $0 \mapsto 0$. We can first define the points p_2 and p_3 on ℓ_2 and ℓ_3 respectively which are closest to $\ell_1(1)$. Then ℓ_1 is on the plane determined by ℓ_2 and ℓ_3 exactly when

$$d(p_2, 0)^2 + d(p_3, 0)^2 = 1.$$

This condition can be written as a formula with the correct zero set, and then the claimed definability follows from properness of $\mathbb{R}^3$.

Together, this all demonstrates the existence of definable sets in correspondence with V_1 and V_2, the projective points and lines respectively, and the definability of the set of incident point-line pairs. To summarize, we can pass to a conservative extension of $(\mathbb{R}^3, d, 0)$ in which we have sorts for V_1, V_2, and the incident pairs in $V_1 \times V_2$, albeit with strange metrics inherited from the weighted sums on coordinate-wise distances. The theory can be extended to ensure that these sorts are interpreted correctly.

From just the incidence structure, we can obtain the real field using classic constructions, for example the following described in chapter VI of [20]. For this result, Veblen and Young require the space satisfy some axioms A (alignment), E (extension), and P (projectivity). The A and E axioms are what one would generally require to call something a projective space, e.g. having unique lines between points, having unique points on pairs of lines, and sufficiently many points and lines so as not to be degenerate. The P axiom is what enables them to verify the field axioms for the multiplication operation constructed below. The spaces we will be interested in are known to be projective spaces over a field and satisfy these axioms, so we will not discuss them in any more detail.

Proposition 8.3.1. *Let P be a projective space satisfying AEP. Let ℓ be a projective line and let 0, 1, and ∞ be distinct points on ℓ. Then $\ell - \{\infty\}$ forms a field with respect to addition and multiplication operations defined in terms of the incidence relation.*

Proof. This appears as Theorem 10 in that chapter VI of [20] where they do the constructions of the operations and check that they satisfy field axioms. That chapter can be seen for a full proof and discussion. We will just describe the constructions of the operations here, since we will make use of the constructions later.

The main concern for us is to emphasize that the construction just consists of repeatedly selecting the point determined by two lines, or the line determined by two points.

We have fixed a line ℓ, and three distinct points 0, 1, and ∞ on ℓ. We can obtain a plane by picking any point p not on ℓ, and taking the union of all points which occur on a line through both p and a point on ℓ. That is, our plane is

$$\bigcup_{p' \in \ell} \{q \in P : q \text{ is on the line through both } p \text{ and } p'\}.$$

We will now assume we are working only in this plane. That is, all points and lines are in this plane.

First we will describe addition. Let ℓ_∞ and ℓ'_∞ be two distinct lines through ∞. Let ℓ_0 be any line through 0. This determines two points:

$$A, \text{ the point where } \ell_0 \text{ meets } \ell_\infty,$$

$$A', \text{ the point where } \ell_0 \text{ meets } \ell'_\infty.$$

For any points x, y on ℓ, we now construct $x + y$ as follows. We obtain two lines:

$$\ell_x\text{, the line between } x \text{ and } A \in \ell_\infty,$$

$$\ell_y\text{, the line between } y \text{ and } B \in \ell'_\infty.$$

From this we obtain two more points and then a line:

$$X\text{, the point where } \ell_x \text{ meets } \ell'_\infty,$$

$$Y\text{, the point where } \ell_y \text{ meets } \ell_\infty,$$

$$\ell_{X+Y}\text{, the line between } X \text{ and } Y.$$

Finally, we get the point representing the sum on our original line:

$$x + y\text{, the point where } \ell_{X+Y} \text{ meets } \ell.$$

Next we describe multiplication. Let ℓ_0, ℓ_1, and ℓ_∞ be any three lines through 0, 1, and ∞ respectively, so that the following points A and B are distinct:

$$A\text{, the point where } \ell_1 \text{ meets } \ell_0,$$

$$B\text{, the point where } \ell_1 \text{ meets } \ell_\infty.$$

For any points x, y on ℓ we construct $x \cdot y$ as follows. We obtain two lines:

$$\ell_x\text{, the line between } x \text{ and } A,$$

$$\ell_y\text{, the line between } y \text{ and } B.$$

From this we get two points and then a line:

$$X\text{, the point where } \ell_x \text{ meets } \ell_\infty,$$

$$Y\text{, the point where } \ell_y \text{ meets } \ell_0,$$

$$\ell_{X \cdot Y}\text{, the line between } X \text{ and } Y.$$

Finally, we get the point representing the product on our original line:

$$x \cdot y\text{, the point where } \ell_{X \cdot Y} \text{ meets } \ell.$$

One can then check that the operations depend on 0, 1, and ∞, but not the other choices, and that the operations satisfy the field axioms. □

Now, to help realign our intuition, we note the following differences in what follows. We will ultimately work in a class of spaces which are not necessarily proper and which have less well-behaved geodesics and flats (e.g. they may branch). We will be focused on special rays starting at the origin rather than the collection of all lines through the origin. Moreover, the incidence structure will not come from these rays relating to 2-flats, but rather from different types of rays and how they relate to each other. For example, there will be type 1 rays representing points, and type 2 rays representing lines. Building theory will let us obtain a projective plane from these configurations of rays.

8.4. $(\mathbb{E}, W)$ **spaces.** In this section we will discuss spaces that admit a special kind of covering by flats. The definition here is based on the axioms in section 4.1.2 of [14] for euclidean buildings, but we have only taken the properties which are also satisfied by symmetric spaces to emphasize this common structure.

The coverings we are interested in will need to be compatible with a special kind of group. We will start by defining these groups and an associated polyhedron. Later, in symmetric spaces and euclidean buildings, this polyhedron will help us quantify how singular or regular a geodesic ray is. This connection will not be clear until then, since it depends on choosing an appropriate $\mathbb{E}$ and W below for the space.

We are differing a bit from the definitions in [14] because of our emphasis on rays at the origin. We focus on these rays because they serve as representatives for $M(\infty)$, the set of all rays modulo being asymptotic, and we cannot easily discuss the relation of being asymptotic in our logic.

For the next few definitions, note that any isometry j of a euclidean space $\mathbb{E}$ induces an isometry $j_\star$ of $\Gamma_\star(\mathbb{E})$ as follows. The idea is essentially to quotient out translations. Let $j \in \mathrm{Isom}(\mathbb{E})$ and let $\gamma \in \Gamma_\star(\mathbb{E})$. Then $j(\gamma)$ is a ray in $\Gamma(\mathbb{E})$ starting at $j(\star)$, which might not be the origin $\star$. However, there is a unique ray in $\Gamma_\star(\mathbb{E})$ which is parallel to $j(\gamma)$, which we can denote $j_\star(\gamma)$. The map $j_\star$ defined by $\gamma \mapsto j_\star(\gamma)$ is an isometry, since angles $\angle_\infty(\gamma, \gamma')$ must be preserved by j and by translation to the parallel rays $j_\star(\gamma)$ and $j_\star(\gamma')$.

Definition 8.4.1. Let $\mathbb{E}$ be a euclidean space $\mathbb{R}^k$ for some $k > 0$, and let W be a subgroup of the isometry group $\mathrm{Isom}(\mathbb{E})$. We define the **spherical part** of W to be the group $W_\star$ of isometries of $\Gamma_\star(\mathbb{E})$ defined by

$$W_\star = \{j_\star : j \in W\}$$

where $j_\star$ is defined as in the preceding paragraph.

Definition 8.4.2. Let $\mathbb{E}$ be a euclidean space, and let W be a subgroup of the isometry group $\mathrm{Isom}(\mathbb{E})$ of $\mathbb{E}$. We say W is an **affine Weyl group** if it is generated by reflections over a hyperplane of codimension 1 and has finite spherical part.

Affine Weyl groups induce a notion of directions in $\mathbb{E}$ modulo W. We formalize this next. Notice that in euclidean spaces, $\angle_\infty$ is just the usual angle between rays. The map given by $\gamma \mapsto \gamma(1)$ between $\Gamma_\star\mathbb{R}^k$ with metric $\angle_\infty$ and the euclidean $(k-1)$-sphere with the angular metric is thus an isometry. Also note that this means the spherical part of an affine Weyl group is always generated by reflections over subspheres of codimension 1.

Definition 8.4.3. Let W be an affine Weyl group on $\mathbb{E}$, and let $W_\star$ be its spherical part. The **anisotropy polyhedron** $\Delta_{(\mathbb{E},W)}$ is the space obtained by taking the quotient of $\Gamma_\star\mathbb{E}$ with metric $\angle_\infty$ modulo the group $W_\star$. We will denote the resulting metric on $\Delta_{(\mathbb{E},W)}$ by d_Δ.

We define a map θ from distinct pairs in $\mathbb{E}$ to $\Delta_{(\mathbb{E},W)}$ as follows. For $p \neq q$ in $\mathbb{E}$, let $\gamma_\star^{p,q}$ be the unique ray starting at $\star$ which is parallel to the unique ray from p through q. We define $\theta(p,q)$ to be the image $\gamma_\star^{p,q}/W_*$ under quotienting by $W_\star$. We call $\theta(p,q)$ the $\Delta_{(\mathbb{E},W)}$-**direction** of (p,q).

It may help to keep in mind that we intend to use the behavior of rays as a proxy for "behavior at infinity". In other words, $W_\star$ is intuitively meant to capture the action induced by W on the "boundary sphere" of $\mathbb{E}$.

Example 8.4.4. Let $\mathbb{E}$ be the euclidean plane $\mathbb{R}^2$, and let W be the group generated by all translations and the reflections through the usual x and y axes. Then W contains reflections through any vertical or horizontal line, but only two of its reflections fix the origin. So, W is an affine Weyl group.

The space $\Gamma_\star\mathbb{E}$ is the set of rays in the plane starting at the origin, and $\angle_\infty$ is just the usual angle between such rays. The group $W_\star$ has order four. One generator of $W_\star$ corresponds to the reflection in W which fixes the origin and reflects through the y-axis. This fixes the two rays pointing along the positive and negative y-axis, and otherwise exchanges rays γ_1 and γ_2 if they are on opposite sides of the y-axis and have the same angle with the positive y-axis. The second generator of $W_\star$ corresponds to the reflection in W which fixes the origin and reflects through the x-axis.

The anisotropy polyhedron in this example can be represented by the quarter plane between and including the positive x and y axes.

A nontrivial example of the $\Delta_{(\mathbb{E},W)}$-direction is the following. Let p be $(-1,0)$ in the standard (x,y) coordinates on the plane, and let q be $(-2,\sqrt{3})$. Then the (unique) ray $\gamma_\star^{p,q}$ starting at $\star$ and parallel to the (unique) ray from p through q is the ray with usual angle $5\pi/6$ counter-clockwise from the positive x-axis. We have that $\theta(p,q) = \gamma_\star^{p,q}/W_\star$, but we will continue and find another representative for this class. We can obtain a representative in the positive quarter-plane by reflecting $\gamma_\star^{p,q}$ over the y-axis (using an element of $W_\star$), to get the ray $\gamma_{\pi/6}$ with usual angle $\pi/6$ counter-clockwise from the positive x-axis. This makes $\theta(p,q)$ the class $\gamma_{\pi/6}/W_\star$ in the anisotropy polyhedron.

The group $W_\star$ does not need to consist of only reflections as in this last example. For example, it may be generated by reflections through the diagonals of a hexagon centered at the origin. In this case $W_\star$ contains some rotations as well.

The reflections in W and $W_\star$ determine certain subsets of $\mathbb{E}$ and $\Gamma_\star(\mathbb{E})$. We will call more attention to these in the section on spherical buildings, but for now we call attention to a few for the affine case.

Definition 8.4.5. Let $\mathbb{E}$ be a euclidean space, and let W be an affine Weyl group on $\mathbb{E}$. A subset of $\Gamma_\star(\mathbb{E})$ is called a **wall** if it is the fixed set of a reflection in $W_\star$. The finitely many connected components of $\Gamma_\star(\mathbb{E})$ with its walls removed are called **chambers**. These walls and chambers determine subsets of $\mathbb{E}$ by taking the union of rays in the wall or chamber, respectively.

Now we will define the class of $(\mathbb{E}, W)$ spaces. As mentioned above, this definition is abstracted from axioms related to building theory, so we will use some of the terminology from that setting. The main idea is that if our space is covered by flats in a way that respects W, then the covering lets us unambiguously extend the map θ to pairs in our space.

Definition 8.4.6. Let M be a CAT(0) space. Let $\mathbb{E}$ be $\mathbb{R}^k$ for some $k > 0$, and let $W \subset \mathrm{Isom}(\mathbb{E})$ be an affine Weyl group.

We say M is an $(\mathbb{E}, W)$ **space** if there is a collection $\mathcal{A}$ of isometric embeddings $\mathbb{E} \to M$ with the properties below. We call $\mathcal{A}$ the **atlas**, the embeddings in $\mathcal{A}$ are called **charts,** the images of these charts are called **apartments**. The image of a wall or chamber of $\mathbb{E}$ under a chart is called a **wall** or **chamber** of M, respectively.

- **(Enough apartments)** Each segment, ray, and line in M is contained in at least one apartment.
- **(W-closed)** $\mathcal{A}$ is closed under precomposition with isometries from W. That is, for all $w \in W$ and $i \in \mathcal{A}$, we have $i \circ w \in \mathcal{A}$.
- **(W-compatible)** Charts in $\mathcal{A}$ are W-compatible in the sense that for any two $i_1, i_2 \in \mathcal{A}$, the partial map $i_1^{-1} \circ i_2$ from part of $\mathbb{E}$ to $\mathbb{E}$ is the restriction of an isometry in W. Phrased differently, if i_1, i_2 have overlapping apartments A_1, A_2 as their images, then there must be $w \in W$ such that i_2 and $i_1 \circ w$ agree on the relevant set $\{p \in \mathbb{E} : i_2(p) \in A_1\}$.

These charts and the map θ on distinct pairs in $\mathbb{E}$ induce a corresponding map we continue to call θ from distinct pairs in M to $\Delta_{(\mathbb{E},W)}$ given by sending $p, q \in M$ to $\theta(i^{-1}(p), i^{-1}(q))$ for any choice of chart $i \in \mathcal{A}$. This is possible because p and q are endpoints of a segment and hence contained in at least one common apartment, and this is well-defined because of W-compatibility. This extended map θ is required to have the following property.

- **(Locally bounded θ)** For any three distinct points $p, q_1, q_2 \in M$, we have
$$d_\Delta(\theta(p, q_1), \theta(p, q_2)) \leq \tilde{\angle}_p(q_1, q_2).$$
That is, the difference between $\Delta_{(\mathbb{E},W)}$-directions for segments starting at the same point p must be bounded by the comparison angle at p.

The following are immediate consequences of the local bound on θ.

Proposition 8.4.7. *Let $x, y, z \in M$ where M is an $(\mathbb{E}, W)$ space.*

(1) If y is on the segment from x to z, then
$$\theta(x, y) = \theta(x, z) = \theta(y, z).$$

(2) If $\tilde{\angle}_x(y, z) = 0$, then
$$\theta(x, y) = \theta(x, z).$$

Moreover, these affirm that any segment, ray, or line $\gamma : I \to M$ can be assigned a $\Delta_{(\mathbb{E},W)}$ direction in a consistent way by choosing any $t_1 < t_2 \in I$ and taking $\theta(\gamma)$ to be $\theta(\gamma(t_1), \gamma(t_2))$.

Definition 8.4.8. If M is an $(\mathbb{E}, W)$ space, then for all geodesics $\gamma : I \to M$ we define $\theta(\gamma)$ to be $\theta(\gamma(t_1), \gamma(t_2))$ where $t_1 < t_2$ are reals in I. We call $\theta(\gamma)$ the $\Delta_{(\mathbb{E},W)}$-direction of γ.

One can also check the following by using the local bound and continuity, and considering the construction of the unique asymptotic ray starting at some p for any ray γ.

Proposition 8.4.9. *Asymptotic rays in M have the same $\Delta_{(\mathbb{E},W)}$-direction.*

8.5. **Symmetric spaces of noncompact type.** We will now discuss symmetric spaces of noncompact type. We will see that they are $(\mathbb{E}, W)$ spaces for a W determined in a natural way from the isometry group. In this case, points in the anisotropy polyhedron correspond to orbits of rays. This is how $\Delta_{(\mathbb{E},W)}$ and the map θ provide a more granular measurement of how singular or regular a ray is. From our perspective this is the key feature of symmetric spaces of noncompact type: the behavior of singular geodesics is accessible and highly structured.

Some good references for this section are chapter II.10 of Bridson and Haefliger's [8], chapter 2.1 of Eberlein's [10], or section 4 of Eberlein's notes [9].

Symmetric spaces provide an important connection between differential geometry and semi-simple Lie groups, and as such they can be approached from several directions. We will use a more restrictive geometric definition below and then discuss the relation to Lie groups. Geometry literature tends to make use of a Riemannian manifold based definition, while the building theory literature tends to use the other approach. First, since symmetric spaces are a special kind of Riemannian manifold, we recall their definition.

Definition 8.5.1. A **Riemannian manifold** is a differentiable manifold M with a collection of scalar products, $\langle \cdot, \cdot \rangle_p$ on the tangent space T_pM for each $p \in M$, such that the scalar products vary continuously with p.

The manifold structure above allows us to define a length for piecewise differentiable paths in M. When M is connected, this can be used to define a distance.

Definition 8.5.2. Let M be a Riemannian manifold and let $c : [a, b] \to M$ be a piecewise differentiable path. The **Riemannian length** of c is given by

$$\int_a^b |c'(t)| dt.$$

When M is connected, we define a metric on M by letting $d(x, y)$ be the infimal Riemannian length among piecewise continuously differentiable paths $c : [0, 1] \to M$ with $c(0) = x$ and $c(1) = y$.

To get a symmetric space, we require our Riemannian manifold to have a special symmetry at every point which is usually thought of as reversing geodesics through that point.

Definition 8.5.3. A **(Riemannian) symmetric space** is a connected Riemannian manifold M such that at each $p \in M$, there is an isometry σ_p of M such that

- $\sigma_p(p) = p$, and
- the differential of σ_p is multiplication by -1 on the tangent space.

The model spaces $\mathcal{M}_\kappa$, which arise as rescaled euclidean spaces, spheres, and hyperbolic spaces are the most familiar examples of symmetric spaces. The symmetries σ_p are indeed given by reversing geodesics: consider the geodesic segment from p to x with initial velocity v, and send x to the point $\sigma_p(x)$ which is $d(p, x)$ far along the segment obtained by leaving p with initial velocity $-v$.

As can be seen from just these examples, symmetric spaces come in a wide variety. We will only be interested in a subclass.

Definition 8.5.4. A symmetric space M is of **noncompact type** if the following hold.

- M is simply connected.
- M is non-positively curved (in the differential geometry sense and hence also CAT(0)).
- M cannot be factored as a Riemannian product $N \times \mathbb{E}$ where $\mathbb{E}$ is a nontrivial euclidean space.

The following two propositions outline the relation between such spaces and semi-simple Lie groups. For proofs or a more in depth discussion, see the first few sections of chapter 2 of [10].

Proposition 8.5.5. *Let M be a symmetric space of noncompact type, and let* $\mathrm{Isom}(M)$ *be the group of isometries on M. The connected component of the identity in* $\mathrm{Isom}(M)$ *is a semi-simple Lie group G with trivial center and no compact factors, and G acts transitively on M.*

Proof. This is 2.1.1 of [10]. The meaning of no compact factors here is based on a decomposition of G into a certain direct product of connected, normal Lie subgroups. $\square$

If we fix a point $p \in M$ then its stabilizer G_p in G is a maximal compact subgroup. One can check that G acts transitively on M (by composing symmetries σ_p). This gives an identification of M with G/G_p. This is how these symmetric spaces arise in general.

Proposition 8.5.6. *Let G be a semi-simple Lie group with trivial center and no compact factors, and let K be a maximal compact subgroup of G. Then $M = G/K$ can be given a G-invariant Riemannian metric, and in this case M is a symmetric space of noncompact type such that the identity component of* $\mathrm{Isom}(M)$ *is G.*

The standard examples for symmetric spaces of noncompact type are the spaces obtained as $\mathrm{SL}_n(\mathbb{R})/\mathrm{SO}_n(\mathbb{R})$ for $n \geq 2$, and more interestingly $n \geq 3$. These examples are standard both because the structure is fairly approachable using just concepts from linear algebra, and because all irreducible symmetric spaces of noncompact type embed nicely into some $\mathrm{SL}_n(\mathbb{R})/\mathrm{SO}_n(\mathbb{R})$. A proof of this can be found in part 2 of the appendix of [9] under the section "The imbedding theorem".

An important feature of a symmetric space is what kinds of flats can be found within it. As a manifold, there is a bound on the dimension of a flat subspace. We give a special name to the maximal dimension that occurs.

Definition 8.5.7. The **rank** of a symmetric space M is the the maximal k for which there are k-flats in M.

Proposition 8.5.8. *Let M be a symmetric space of noncompact type and rank k, and let G be the identity component of* $\mathrm{Isom}(M)$.

(1) *If F_1 and F_2 are k-flats in M, and if $p_1 \in F_1$ and $p_2 \in F_2$, then there is $g \in G$ such that $g(p_1) = p_2$ and $g(F_1) = F_2$.*
(2) *If γ is a geodesic line in M, then there is a maximal flat containing γ.*

Proof. See appendix part 2, Proposition 26 of [9], or section 2.10 of [10]. □

The transitivity of G on the k-flats lets us define an atlas and view M as an $(\mathbb{E}, W)$ space. This construction is outlined in section 5.2 of Kleiner and Leeb's [14], and we summarize it now.

Definition 8.5.9. Let M be a symmetric space of noncompact type and rank k. Let G be the identity component of $\mathrm{Isom}(M)$. We define the **maximal** $(\mathbb{E}, W)$ **atlas** $\mathcal{A}$ on M as follows.

(1) Let $\mathbb{E}$ be $\mathbb{R}^k$.
(2) Choose any k-flat F in M. Let $S_G(F)$ be the setwise stabilizer of F in G, and let $P_G(F)$ be the pointwise stabilizer of F in G. Define
$$W' = S_G(F)/P_G(F).$$
Note that W' can be viewed as a subgroup of $\mathrm{Isom}(F)$.
(3) $W \subset \mathrm{Isom}(\mathbb{E})$ is defined using W' and an isometry $i : \mathbb{E} \to F$. That is,
$$W = \{i^{-1} \circ w' \circ i : w' \in W'\}.$$
(4) Let $\mathcal{A}$ consist of all isometric embeddings $\mathbb{E} \to M$ such that W is preserved, in the sense that if $i \in \mathcal{A}$, then
$$i \circ W \circ i^{-1}$$
is exactly the group
$$S_G(i(\mathbb{E}))/P_G(i(\mathbb{E})).$$

Kleiner and Leeb note that this W is generated by reflections and is an affine Weyl group on $\mathbb{E}$, and the transitivity of G on k-flats gives the following.

Proposition 8.5.10. *M is an $(\mathbb{E}, W)$ space with the maximal $(\mathbb{E}, W)$ atlas described above. Every maximal flat in M is the image of some chart in this atlas.*

Discussions of these spaces often work out the example $\mathrm{SL}_n(\mathbb{R})/\mathrm{SO}_n(\mathbb{R})$ in extensive detail, as can be seen in the sources we have mentioned. They show that this space is isometric to $P = \mathrm{P}_1(n, \mathbb{R})$, the space of positive definite symmetric matrices with determinant 1, carrying a certain Riemannian metric. The maximal flats in P are dimension $n-1$ and are conjugates of a canonical flat given by the diagonal matrices in P. The group $W_\star$ in this case corresponds to permutations of the coordinates in these diagonal matrices. In particular, when $n = 3$, the group $W_\star$ is isomorphic to the permutation group on 3 elements, and 3 of the isometries in $W_\star$ are reflections of $\mathbb{E} = \mathbb{R}^2$. The anisotropy polyhedron is isometric to a circular arc of length $2\pi/6$, including the endpoints. The endpoints of this arc correspond to two distinct orbits of singular rays, and the interior points correspond to distinct orbits of regular rays.

8.6. **Euclidean buildings.** In this section we discuss euclidean buildings, which are $(\mathbb{E}, W)$ spaces with an additional rigidity property relating initial angles $\angle_p$ to angles at infinity $\angle_\infty$. The terminology in the literature is somewhat confusing when it comes to the names of these buildings. There are various definitions of discrete and non-discrete euclidean buildings, affine buildings, $\mathbb{R}$-buildings, and more general Λ-buildings depending on the generality and purpose of the discussion.

We are essentially following the definition of euclidean buildings given in Kleiner and Leeb's [14]. But, since we will need to work primarily with rays at the origin in our setting, we have changed all of the definitions to reflect this and also to fit our notation. Again, there is no substantial difference since $\Gamma_\star\mathbb{E}$ is a set of representatives for those classes in $\mathbb{E}(\infty)$, the set of rays modulo being asymptotic. For proofs of the equivalence of Kleiner and Leeb's definition with other definitions, readers are usually referred to II.2.7 of Parreau's thesis [17], but [6] would also be useful.

First, we will name the set of possible angles at infinity which can occur between rays with given $\Delta_{(\mathbb{E},W)}$-directions. Note that this set is finite, since $W_\star$ is finite by definition.

Definition 8.6.1. Let W be an affine Weyl group on $\mathbb{E}$. For $\alpha_1, \alpha_2 \in \Delta_{(\mathbb{E},W)}$, we define $D(\alpha_1, \alpha_2)$ to be the finite set of possible distances $\angle_\infty(\gamma_1, \gamma_2)$ between elements $\gamma_1, \gamma_2 \in \Gamma_\star\mathbb{E}$ such that $\alpha_1 = \gamma_1/W_\star$ and $\alpha_2 = \gamma_2/W_\star$.

Example 8.6.2. Let $\mathbb{E} = \mathbb{R}^2$ and let W be such that $W_\star$ is generated by reflecting rays over the x-axis and y-axis. In the earlier example 8.4.4, we saw that $\Delta_{(\mathbb{E},W)}$ can be matched with the rays in the positive quarter-plane. Consider the classes $\alpha_0, \alpha_{\pi/6} \in \Delta_{(\mathbb{E},W)}$ represented by the ray pointing along the positive x-axis and the ray with angle $\pi/6$ from the positive x-axis, respectively. The only other member of α_0 is the ray with angle π from the positive x-axis. The other members of $\alpha_{\pi/6}$ are the rays with angles $-\pi/6$, $\pi-(\pi/6)$, and $\pi+(\pi/6)$ from the positive x-axis. The way that these contribute members to $D = D(\alpha_0, \alpha_{\pi/6})$ is described in the following table, which matches pairs of rays in $\alpha_0 \times \alpha_{\pi/6}$ with their distance. Remember that the metric $\angle_\infty$ on $\Gamma_\star\mathbb{E}$ is the unsigned angle between rays. The angles which occur are redundant, but we obtain two distinct elements of D, not just the original $\pi/6$ we start with.

angle of ray in α_0	angle of ray in $\alpha_{\pi/6}$	element of D
0	$\pi/6$	$\pi/6$
0	$-\pi/6$	$\pi/6$
0	$\pi+(\pi/6)$	$\pi-(\pi/6)$
0	$\pi-(\pi/6)$	$\pi-(\pi/6)$
π	$\pi/6$	$\pi-(\pi/6)$
π	$-\pi/6$	$\pi-(\pi/6)$
π	$\pi+(\pi/6)$	$\pi/6$
π	$\pi-(\pi/6)$	$\pi/6$

So in this case, $D = \{\pi/6, \pi-(\pi/6)\}$. Choosing α away from the boundar of $\Delta_{(\mathbb{E},W)}$ would yield a 4-element set $D(\alpha, \alpha_{\pi/6})$. For example, even $D(\alpha_{\pi/6}, \alpha_{\pi/6})$ would consist of 4 distinct elements.

Definition 8.6.3. A **euclidean building** is an $(\mathbb{E}, W)$ space whose map θ has the following additional property.

- **(Angle rigidity)** For any three distinct $p, q_1, q_2 \in M$, the initial angle $\angle_p(q_1, q_2)$ is in the finite set $D(\theta(p, q_1), \theta(p, q_2))$.

This additional property of euclidean buildings lets us prove the following result about the behavior of rays starting at the same point. We will make use of this next proposition when proving definability results later.

Proposition 8.6.4. ***(Rays diverge piecewise linearly)*** *Let M be a euclidean building where the affine Weyl group W is transitive on points in $\mathbb{E}$. Let γ_1, γ_2 be two rays starting at the same point $p \in M$. The function $F : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ given by*

$$F(t) = d(\gamma_1(t), \gamma_2(t))$$

is nondecreasing, piecewise linear, and each linear part has slope $2 \cdot \sin(\alpha/2)$ for some $\alpha \in D(\theta(\gamma_1), \theta(\gamma_2))$.

Proof. Before we begin, we make the following note. In Lemma 4.1.2 of [14], they use angle rigidity directly to show that γ_1 and γ_2 must either initially coincide or else initially span a flat triangle and have an initial angle in $D(\theta(\gamma_1), \theta(\gamma_2))$. Consequently, in a neighborhood of 0, the function $F(t)$ and the claimed property of its slope are just consequences of the geometry of triangles in the euclidean plane. The proposition we are currently considering can be thought of as a generalization saying that this initial nice behavior of F is just the first piece of a piecewise property.

Also, note that knowing this initial behavior of F means in particular that it is initially nondecreasing. The argument makes use of the fact that in CAT(0) spaces, this function F is convex (prop 5.3.6). Knowing that F is initially nondecreasing then implies that F must be nondecreasing everywhere.

In Theorem 3.3 of [6], among many other equivalences it is shown that these buildings satisfy a "large atlas" condition they abbreviate as (LA). Informally, this says that any two chambers are initially contained in a common apartment. More precisely, for any pair of affine Weyl chambers C_1 and C_2, there are neighborhoods $C_1' \subseteq C_1$ and $C_2' \subseteq C_2$ of the tips in C_1 and C_2 respectively for which there is a common apartment containing both C_1' and C_2'.

Consider any $t_0 \in [0, \infty)$. Because M is an $(\mathbb{E}, W)$ space, the ray $\gamma_1([t_0, \infty))$ is contained in some apartment A_{1,t_0}, and the ray $\gamma_2([t_0, \infty))$ is contained in some apartment A_{2,t_0}. Since we have assumed W is transitive on points, we can take A_{1,t_0} and A_{2,t_0} so that these rays are in chambers $C_1 \subseteq A_{1,t_0}$ and $C_2 \subseteq A_{2,t_0}$ with tips $\gamma_1(t_0)$ and $\gamma_2(t_0)$ respectively. By the above paragraph, there are C_1', C_2', and A such that

- C_1' contains an initial segment of $\gamma_1([t_0, \infty))$,
- C_2' contains an initial segment of $\gamma_2([t_0, \infty))$,
- A is an apartment containing both C_1' and C_2'.

So, there must be some $t_1 > t_0$ such that the segments $\gamma_1([t_0, t_1])$ and $\gamma_2([t_0, t_1])$ are both contained in A. Thus, for $t \in [t_0, t_1]$, the function F is a function of distances between segments in a euclidean plane. This reduces the claim to euclidean geometry of line segments. The monotonicity of F justifies computing the slope of $F([t_0, t_1])$ by translating the segment $\gamma_1([t_0, t_1])$ to a parallel segment γ_1' also having $\theta(\gamma_1') = \theta(\gamma_1)$ but with $\gamma_1'(t_0) = \gamma_2(t_0)$. It follows from angle rigidity and the geometry of triangles in the euclidean plane that the slope is as claimed. □

Corollary 8.6.5. *Let M, W, γ_1, γ_2, and F be as above. There is a finite $N \in \mathbb{N}$ which depends on W but not on γ_1 or γ_2 such that F consists of at most N many linear segments.*

Proof. There is a uniform bound N on the cardinality of $D(\theta(\gamma_1), \theta(\gamma_2))$ independent of γ_1, γ_2 because the number of reflections through the origin in W is finite. The claim then follows from the properties of F given in the previous proposition. □

These preceding claims are useful in our logic because they show that we can know something about the asymptotic behavior of γ_1 and γ_2 by considering finitely many pairs $\gamma_1(t), \gamma_2(t)$.

Example 8.6.6. One of the simplest nontrivial examples of a euclidean building is an $\mathbb{R}$-tree. Such a space M can be constructed by starting with a copy of $\mathbb{R}$, inductively adjoining a new copy of $(\mathbb{R}, 0)$ at every point, and repeating this process. The shortest path is used to determine the metric when new copies of $\mathbb{R}$ are added. That is, whenever a new copy $(R, 0')$ of $(\mathbb{R}, 0)$ is added to M at a point p, so that $0'$ and p are identified, the metric is extended so that for old points $x \in M$ and new points $y \in R$, we have

$$d(x, y) = d(x, p) + d(p, y).$$

This yields a CAT(0) space which is a tree-like object that "branches everywhere" in the sense that at every point p, there are $2^{\aleph_0}$ many distinct segments leaving p. The construction ensures that there is a unique geodesic segment between each pair of points, and that every segment extends to a full geodesic line in infinitely many ways.

For an $(\mathbb{E}, W)$ space structure on M, we take $\mathbb{E} = \mathbb{R}$ and let W be the full isometry group of $\mathbb{R}$, including all translations and reflections. The space $\Gamma_\star\mathbb{E}$ consists of only two rays: one leaving $0 \in \mathbb{R}$ in each direction. The group $W_\star$ has one nontrivial element exchanging these two rays, and $\Delta_{(\mathbb{E},W)}$ is just a singleton since the two rays collapse.

For our atlas, we can take all isometries $\mathbb{R} \to M$. Note that this includes more copies of $\mathbb{R}$ than just those used in the construction of M. Lines are trivially contained in apartments, since in this case the apartments are exactly the lines. Segments and rays extend to lines and hence are contained in apartments. The angle rigidity property is satisfied because all rays have the same $\Delta_{(\mathbb{E},W)}$-direction by virtue of $\Delta_{(\mathbb{E},W)}$ being a singleton and any pair of rays leaving the same point in M either have initial angle 0 or π. Any pair which does not coincide forever must have a point where they begin diverging with angle π, demonstrating the claims about $d(\gamma_1(t), \gamma_2(t))$.

Trees are rather degenerate cases of euclidean buildings because the maximal flats are dimension 1. More interesting cases involve higher dimension flats. For example, we could construct an analog of the above example by starting with a plane and inductively attaching new planes by gluing along lines. Depending on which lines are chosen for gluing or what atlas is intended, this can lead to $(\mathbb{E}, W)$ space structures where $\mathbb{E} = \mathbb{R}^2$ but W does not contain all reflections. In such cases, $\Delta_{(\mathbb{E},W)}$ is a nondegenerate polytope.

8.7. **Spherical buildings.** In this section, we define a CAT(1) analog of the buildings we just discussed. Jacques Tits developed these spherical buildings before developing affine buildings and the later generalizations that would lead to the euclidean buildings defined above. The original purpose of these objects was to build a combinatorial, geometric object to help classify algebraic groups.

The order in which we have defined things and the approach differs because we will work with spherical buildings as structures found on $\Gamma_\star M$ for both symmetric spaces and euclidean buildings. That connection will be discussed in the next section. For now, we just define spherical buildings as objects on their own. We will again follow the development in chapter 3 of Kleiner and Leeb's [14] but with minor adaptations for our setting.

The first definition below is analogous to the definitions of $\Gamma_\star\mathbb{E}$ and $W_\star$ in the euclidean building sections. Spherical buildings will require defining an atlas of isometries from a euclidean k-sphere, rather than from some $\mathbb{R}^k$.

Definition 8.7.1. Let $\mathbb{S}$ be the unit sphere of some euclidean space, equipped with the angular metric $\angle$. Let W be a subgroup of the isometry group $\mathrm{Isom}(\mathbb{S})$. We say W is a **Weyl group** if it is finite and generated by reflections over a subsphere of codimension 1.

In this setting we also get an anisotropy polyhedron and quotient map.

Definition 8.7.2. The **anisotropy polyhedron** $\Delta_{(\mathbb{S},W)}$ is the space obtained by taking the quotient of $\mathbb{S}$ by the group W. We will denote the resulting metric by d_Δ. We define θ to be the quotient map from $\mathbb{S}$ to $\Delta_{(\mathbb{S},W)}$. We call $\theta(p)$ the $\Delta_{(\mathbb{S},W)}$**-direction** of p.

Moreover, since W is finite, we again have a finite set of possible distances between preimages of $\Delta_{(\mathbb{S},W)}$-directions.

Definition 8.7.3. For $\alpha_1, \alpha_2 \in \Delta_{(\mathbb{S},W)}$, we write $D(\alpha_1, \alpha_2)$ to denote the finite set of possible distances $\angle(p, q)$ which occur between elements $p, q \in \mathbb{S}$ such that $\alpha_1 = p/W$ and $\alpha_2 = q/W$.

We now define spherical buildings as CAT(1) spaces together with a W-respecting atlas of embeddings of $\mathbb{S}$. Again, the key observation is that such an atlas lets us unambiguously extend the map θ to the space M.

Definition 8.7.4. Let M be a CAT(1) space. Let $\mathbb{S}$ be a euclidean k-sphere, and let $W \subset \mathrm{Isom}(\mathbb{S})$ be a Weyl group. We say M is a **spherical building** if there is a collection $\mathcal{A}$ of isometric embeddings $\mathbb{S} \to M$ with the properties below. We call $\mathcal{A}$ the **atlas**, the embeddings in $\mathcal{A}$ are called **charts,** the images of these charts are called **apartments**.

- **(Enough apartments)** For any pair of points $p, q \in M$ there is at least one apartment containing both p and q.
- **(W-closed)** $\mathcal{A}$ is closed under precomposition with isometries from W. That is, for all $w \in W$ and $i \in \mathcal{A}$, we have $i \circ w \in \mathcal{A}$.
- **(W-compatible)** Charts in $\mathcal{A}$ are W-compatible in the sense that for any two $i_1, i_2 \in \mathcal{A}$, the partial map $i_1^{-1} \circ i_2$ from part of $\mathbb{S}$ to $\mathbb{S}$ is the restriction of an isometry in W. Phrased differently, if i_1, i_2 have overlapping apartments A_1, A_2, then there must be $w \in W$ such that i_2 and $i_1 \circ w$ agree on the relevant set $\{p \in \mathbb{S} : i_2(p) \in A_1\}$.

These charts and the map θ on points in $\mathbb{S}$ induce a corresponding map we continue to call θ from points in M to $\Delta_{(\mathbb{S},W)}$ given by sending $p \in M$ to $\theta(i^{-1}(p))$ for any choice of chart $i \in \mathcal{A}$. This is well-defined because of W-compatibility.

The axioms easily imply that the map θ induced on M above satisfies a discreteness condition which is the analog of the angle rigidity in euclidean buildings.

Proposition 8.7.5. *For points p, q in a spherical building M, we always have that $d(p, q)$ is in the finite set $D(\theta(p), \theta(q))$.*

The atlas on a spherical building determines a simplicial complex. We summarize some of the terminology and simple observations.

Definition 8.7.6. The set of fixed points (in $\mathbb{S}$) of a given reflection from W is called the **wall** of that isometry.

The images of walls under charts are called the **walls** of the building.

Proposition 8.7.7. *Each wall is a subsphere whose dimension is one lower than that of $\mathbb{S}$.*

Proof. Reflections of a sphere have codimension 1 subspheres as their fixed point sets. □

The collection of walls in $\mathbb{S}$ is finite because W is finite. We use the walls to identify other important subsets of $\mathbb{S}$. The points in $\mathbb{S}$ not lying in a wall are divided into pairwise isometric, open, convex sets.

Definition 8.7.8. Let $\mathbb{S}'$ denote the set of points $p \in \mathbb{S}$ which are not in any wall. The closures of the convex components of $\mathbb{S}'$ are called **chamber**s.

The images of chambers under charts are called the **chambers** of the building.

Proposition 8.7.9. *Any chamber in $\mathbb{S}$ is a fundamental domain for W and can also be viewed as a finite intersection of closed hemispheres in $\mathbb{S}$.*

Proof. The chambers arise as connected components separated by walls. So reflections in W exchange adjacent chambers, and any one can be sent to another by a finite sequence of such reflections. To see the hemisphere part of the claim, notice that each chamber C intersects only finitely many walls. Since each such wall is a reflection, it determines a subsphere of codimension 1 which divides $\mathbb{S}$ into two hemispheres, one of which contains C. □

Corollary 8.7.10. *Any chamber in the spherical building is a fundamental domain for the building under isometries given by composing charts and can also be viewed as a finite intersection of images of closed hemispheres under charts (half-apartments).*

Proof. This follows from the definition of charts and the W-closure and compatibility of the atlas. □

We get lower dimensional simplices by intersecting chambers with walls.

Definition 8.7.11. In a spherical building, the intersection of a chamber with one or more walls is called a **face** of that chamber. We also call a face a **vertex** if it is a single point.

So we have described the simplicial structure of a building. The top level simplices are given by the chambers, and the faces give the lower simplices.

The chambers determine an incidence relation on the points in the building. In particular, this restricts to an incidence relation on the collection of vertices.

Definition 8.7.12. Let M be a spherical building, and let $p_1, p_2 \in M$. We say p_1 and p_2 are **incident** if there is a chamber containing both.

For example, if $\mathbb{S}$ is a circle and its chambers are arcs, then the vertices of a $(\mathbb{S}, W)$ spherical building will be the images under charts of endpoints of those arcs. Two such vertices are incident when they are joined by the image of an arc under some chart.

The picture to have in mind for spherical buildings is that they look like a collection of copies of $\mathbb{S}$ (the apartments) joined together along walls.

Example 8.7.13. Let M be the space formed by taking two copies of the 1-sphere (i.e. a circle) and joining them north pole to north pole and south pole to south pole. Use the usual angular metric within each circle, and use the length of the shortest path for distances between points in opposite circles.

Let $\mathbb{S}$ be the 1-sphere, and let W be the group generated by the two following reflections of $\mathbb{S}$:

- the reflection exchanging the north and south poles,
- the reflection fixing these poles and exchanging the midpoints between them.

For the atlas on M, take the four embeddings of $\mathbb{S}$ which send the north pole of $\mathbb{S}$ to one of the two intersection points in M. This makes M a spherical building. The anisotropy polyhedron $\Delta_{(\mathbb{S},W)}$ is a closed quarter-circle of $\mathbb{S}$.

Proposition 8.7.14. *Let M be a spherical building. Any geodesic in M is contained in some apartment. Any isometrically embedded unit sphere in M is contained in an apartment.*

Proof. This is Corollary 3.9.2 in [14]. It is a consequence of a more general result showing that convex subsets of M which are isometric to parts of a euclidean sphere must be contained in an apartment. □

8.8. **Spherical buildings at infinity.** We have seen that both symmetric spaces and euclidean buildings M are $(\mathbb{E}, W)$ spaces, that CAT(0) spaces always have a CAT(1) metric $\angle_\infty$ on $\Gamma_\star M$, and that $\Gamma_\star \mathbb{E}$ is isometric to a euclidean sphere. This yields a natural candidate atlas for $\Gamma_\star M$, as in the definition of a spherical building using the sphere $\Gamma_\star \mathbb{E}$ and the group $W_\star$. The gist of the process is to note that each copy of $\Gamma_\star \mathbb{E}$ in $\Gamma_\star M$ is a sphere $\mathbb{S}$, and induce a spherical atlas from the euclidean atlas on M.

Checking that the atlas obtained on $\Gamma_\star M$ is actually a spherical building mostly comes down to checking the "enough apartments" axiom. In general, the spherical atlas obtained from an $(\mathbb{E}, W)$ space fails to have this property, as in the next example. This shows how the angle rigidity of euclidean buildings and the asymptotic behavior of angles in symmetric spaces plays an important role.

Example 8.8.1. Let $M = \mathbb{R}^2$. We will view M as an $(\mathbb{E}, W)$ space where $\mathbb{E} = \mathbb{R}$ and W is the full isometry group on $\mathbb{R}$. Let the atlas be the set of all embeddings $\mathbb{R} \to M$. We have that $\Gamma_\star\mathbb{E}$ is the unit sphere S^0 in $\mathbb{R}$, and $W_\star$ contains the isometry exchanging the two points of S^0. Thus $\Delta_{(\mathbb{E},W)}$ can be identified with a single point. Note that M with this atlas fails the angle rigidity property of euclidean buildings; almost all rays starting at the same point have an initial angle which is neither 0 nor π.

Consider the collection of embeddings $S^0 \to \Gamma_\star M$ induced by the embeddings $\mathbb{R} \to M$. Let $\gamma_1, \gamma_2 \in \Gamma_\star M$ be such that $\angle_\infty(\gamma_1, \gamma_2) = \pi/2$. There is no embedding $S^0 \to \Gamma_\star M$ which contains both γ_1 and γ_2 in its image.

We should make sure our intuition is accurate and clarify some of the overloaded notation. The $\star$ in $\Gamma_\star\mathbb{E}$ refers to the basepoint of some $\mathbb{R}^k$, while the $\star$ in $\Gamma_\star M$ refers to the basepoint of M. Because the relevant metric is $\angle_\infty$, an isometric copy of $\Gamma_\star\mathbb{E}$ in $\Gamma_\star M$ is not necessarily of the form $\Gamma_\star F$ for a flat F through the basepoint $\star$ in M. That is, there are euclidean spheres appearing in $\Gamma_\star M$ which do not arise as the set of rays in a flat through $\star$ in M. This is perhaps best explained by an example.

Example 8.8.2. Recall the $\mathbb{R}$-tree that we saw in example 8.6.6. Consider the basepoint $\star$ and a point p distinct from $\star$.

Let γ_1 and γ_2 be two rays starting at $\star$ which pass through p but then diverge from each other (i.e. take different branches at p). Then $\angle_\infty(\gamma_1, \gamma_2) = \pi$, so $\{\gamma_1, \gamma_2\}$ is isometric to the unit 0-sphere in $\mathbb{R}^1$. But $\gamma_1 \cup \gamma_2$ is not a flat in M, it is the union of a flat through p and an additional segment from p to $\star$. There is no flat through $\star \in M$ which contains the entire lengths of both γ_1 and γ_2.

This situation has to do with our insistence on primacy of rays at the origin as the representatives of $M(\infty)$. Of course, the rays from p on, i.e. $\gamma_1([d(\star, p), \infty))$ and $\gamma_2([d(\star, p), \infty))$, are also rays and their union is a flat. These tail rays are asymptotic to the original γ_1 and γ_2 respectively, and so remain in the same classes in $M(\infty)$.

For contrast, consider two rays γ_3 and γ_4 which start at $\star$ and diverge immediately. Then $\{\gamma_3, \gamma_4\}$ is also a unit 0-sphere. But in this case, we can view $\{\gamma_3, \gamma_4\}$ as the rays in a flat through $\star$, namely their union.

In any case, our main interest is that $\Gamma_\star M$ has a spherical building structure strongly related to the $(\mathbb{E}, W)$ space structure on M.

Proposition 8.8.3. *Let $(M, d, \star)$ be a symmetric space of noncompact type or a euclidean building, and let $\mathcal{A}_M$ be the maximal $(\mathbb{E}, W)$ atlas on M. Then $\Gamma_\star M$ with metric $\angle_\infty$ is a spherical building based on $(\Gamma_\star\mathbb{E}, W_\star)$ whose atlas consists of the embeddings $\Gamma_\star\mathbb{E} \to \Gamma_\star M$ induced by charts $\mathbb{E} \to M$ in $\mathcal{A}_M$.*

Proof. A proof for the symmetric space makes up section I.2 of [7] or can be found as appendix 5 (culminating in section 7 thereof) in [2]. The example $\mathrm{SL}_n(\mathbb{R})/\mathrm{SO}_n(\mathbb{R})$ is also in II.10 of [8].

A proof for the euclidean building case is more straightforward because of the similarity in definitions, and can be found as Proposition 4.2.1 in [14].

We will briefly comment on the proofs. Because of the way we have defined things, this is almost purely definition chasing. The questionable aspect is the "enough apartments" property for the spherical building. For this, one needs to know that any pair of elements from $\Gamma_\star M$ are both contained in at least one subset of $\Gamma_\star M$ isometric to the sphere $\Gamma_\star \mathbb{E}$. Our comment and example above demonstrate that this is not the same as checking that every pair of rays at $\star$ in M is contained in a flat.

For the euclidean building case, Kleiner & Leeb prove some helping lemmas and use the angle rigidity property to complete the task. Roughly, they first show that if $\angle_\infty(\gamma_1, \gamma_2) = \pi$, then the angle rigidity and monotonicity of initial angles as we move out along a rays lets us find a geodesic line whose positive direction coincides with a tail of γ_1 and whose negative direction coincides with a tail of γ_2. By the "enough apartments" axiom for euclidean buildings, this geodesic line is contained in a flat (which necessarily corresponds to a copy of $\Gamma_\star \mathbb{E}$), and so the claim is verified for the case of two rays with π as their angle at infinity. They then handle the cases where $\angle_\infty(\gamma_1, \gamma_2) < \pi$ using that case and some convexity properties of spheres in $\Gamma_\star M$.

For the symmetric space case, things are less direct. The intuition suggested by Kleiner and Leeb here is that while symmetric spaces do not satisfy the angle rigidity property enjoyed by euclidean buildings, pairs of rays in symmetric spaces tend toward the rigid angles as we travel along them toward infinity. While two rays might not occupy a common flat, there is always at least one flat which they will both be asymptotic to. Typically, one would need to pass to other representatives of the classes of γ_1 and γ_2 in $M(\infty)$ in order to get rays that actually intersect that flat. We will look at an example in the hyperbolic plane below to demonstrate this. Verifying the "enough apartments" axiom for spherical buildings requires checking that this behavior is actually happening in symmetric spaces. To do this, one has to understand the asymptotic behavior of rays well. In the sources mentioned, this is accomplished by studying the Iwasawa decomposition of the Lie group corresponding to the symmetric space, and understanding the action of certain group elements on the classes of rays making up $M(\infty)$. □

Example 8.8.4. Consider the open disk model of the hyperbolic plane M, which is a symmetric space of noncompact type. Here, lines of M are given by diameters of the disk as well as circular arcs which meet the boundary of the disk at right angles. Consider two rays γ_1 and γ_2 starting at the origin with initial angle $\pi/2$. This pair has $\angle_\infty(\gamma_1, \gamma_2) = \pi$, and $\{\gamma_1, \gamma_2\}$ is isometric (using $\angle_\infty$) to a unit 0-sphere.

We will now locate the corresponding 1-flat. Let b_1 and b_2 denote the boundary points of the disk toward which γ_1 and γ_2 tend, respectively. Let F be the circular arc in the disk meeting b_1 and b_2 at right angles. Then F is the 1-flat we are looking for. Note that F does not intersect the rays γ_1, γ_2 at all in M.

Consider the pairs of rays that can be obtained by selecting t very large and looking at rays from $\gamma_1(t)$ toward b_1 and from $\gamma_1(t)$ toward b_2. For large values of t, the unions of such rays approach F in the sense that points on these rays can be made arbitrarily, uniformly close to F in the distance d on M.

Finally, we will identify a pair of rays in F which are asymptotic to γ_1 and γ_2. Let p be any point in F, and let γ_1', γ_2' be rays in F starting at p and heading toward b_1 and b_2 respectively. The rays γ_1' and γ_2' are then asymptotic to γ_1 and γ_2 respectively.

In the remainder of this section, we note the relation between the spherical building just obtained for a symmetric space and the building classically associated to symmetric spaces via the more Lie theoretic approach. The group $W_\star$ above is isomorphic to the Weyl group of the symmetric space usually defined algebraically. The relevant connection between rays $\Gamma_\star M$, the equivalence classes of rays $M(\infty)$, and certain subgroups of the Lie group G associated to M is explained in [7] for example. We will just include one relevant proposition below.

Recall that the isometries in G induce isometries on the set of rays and hence on $M(\infty)$ with the metric $\angle_\infty$.

Definition 8.8.5. A **Borel subgroup** of G is a maximal connected solvable algebraic subgroup of G. A **parabolic subgroup** of G is a subgroup $P \leq G$ which contains a Borel subgroup.

Proposition 8.8.6. *Let $M = G/K$ be a symmetric space of noncompact type.*

(1) *For each point $\bar\gamma \in M(\infty)$, the stabilizer of $\bar\gamma$ in G is a parabolic subgroup.*
(2) *Every proper parabolic subgroup of G is the stabilizer of some point in $M(\infty)$.*

Proof. See I.2.6 of [7] or the example $\mathrm{SL}_n(\mathbb{R})/\mathrm{SO}_n(\mathbb{R})$ explained in II.10 of [8]. □

Thus, $\Gamma_\star M$ can be thought of as parametrizing the parabolic subgroups of G. The algebraic definition of the spherical building $\Delta(G)$ associated to G is typically given as a simplicial complex built from the proper parabolic subgroups of G. The simplices are Δ_P for each such $P \subset G$. These Δ_P are vertices when P is maximal, and we also have that Δ_P contains Δ_Q as a face exactly when $P \subset Q$ in G. The chambers correspond to minimal P (i.e. the Borel subgroups). Because of the correspondence between rays and parabolic subgroups above, the algebraic definition is consistent with the geometric definition of the spherical building we are using.

8.9. **Projective planes in spherical buildings.** In general, a spherical building might not have a projective plane associated to it. For a trivial example, if the building is based on $\mathbb{S}$ being the 0-sphere, then the incidence relation is uninteresting because the chambers are degenerate. For higher rank buildings we can extract a meaningful incidence geometry, but in general it may still fail to satisfy strong enough axioms to construct a field. The geometric axioms required to get associativity or commutativity may fail, for example.

We will only be concerned with the buildings arising from symmetric spaces of noncompact type, with rank $k \geq 2$ and corresponding to G/K where G is absolutely simple and defined over $\mathbb{R}$. In this case, our spherical building at infinity is a geometric realization of the building constructed from parabolic subgroups of $G(\mathbb{R})$ and determines the field $\mathbb{R}$, and the spherical building of the $\mathcal{U}$ asymptotic cone is a geometric realization of the spherical building constructed from parabolic subgroups of $G({}^{\rho}\mathbb{R}_{\mathcal{U}})$, which we discuss near the end of the thesis.

The key to extracting the projective plane, when it is possible, is to look at the incidence relation between certain subsets of vertices in the spherical building. We first impose a labeling of the vertices in the anisotropy polyhedron, which then extends to a labeling of vertices in the entire building by looking at preimages under θ. In other words, we label the vertices of chamber and then extend the labeling using orbits under W. These labelings come from the Dynkin or Coxeter diagram of our group W, which are tools used in the classification of semisimple Lie algebras. We will define the Dynkin diagram here.

Definition 8.9.1. Let W be a Weyl group for a sphere $\mathbb{S}$, and let R be a minimal generating set of reflections for W. The **Dynkin diagram** D for W is a graph with weighted edges defined as follows. The nodes of D are the generating reflections R. Each pair of nodes w and w' are connected by an edge with weight $|w \cdot w'| - 2$, where $|w \cdot w'|$ is the order of $w \cdot w'$ in W. We will consider two nodes to be **adjacent** if they are connected by an edge with a positive weight.

Because the vertices of a chamber each correspond to an intersection of $k-1$ walls (i.e. fixed sets of reflections), each vertex in a chamber can be associated with exactly one reflection and hence one of the nodes in the Dynkin diagram. The finite Dynkin diagrams have been classified, and the buildings in our cases have diagrams containing the following type of Dynkin diagram.

Definition 8.9.2. A Dynkin diagram is of type A_n if it consists of n nodes $v_1, \ldots, v_n$, where there is an edge of weight 1 between each pair v_i, v_{i+1}, and an edge of weight 0 between all other pairs.

Such diagrams look like a line of nodes joined by single edges. For our rank k symmetric spaces with the maximal atlas, our spherical buildings at infinity have chambers with k many vertices each, corresponding to the k distinct orbits of maximally singular rays. The Dynkin diagram similarly has k many nodes. We label the vertices of the anisotropy polyhedron $1, \ldots, k$ using the correspondence with nodes in the Dynkin diagram. We extend this labeling to the spherical building using θ^{-1} and then to $\Gamma_\star(M)$ using the atlas. The set of label 1 and 2 vertices, together with the incidence relation defined by being contained in a common chamber, is then isomorphic to a projective plane over $\mathbb{R}$. So, we can use the constructions we described in Proposition 8.3.1. This is how we will build a projective plane from our rays.

From here on, we will assume such a labeling is fixed for each of our buildings. An exposition can be found in II.10 of [8] or 2.13.8 of [10] of the isomorphism between the building at infinity for $\mathrm{SL}_n(\mathbb{R})/\mathrm{SO}_n(\mathbb{R})$ and the usual projective space constructed from linear subspaces of $\mathbb{R}^n$. A more general discussion of buildings and the relationship with the fundamental theorem of projective geometry can be found in the survey paper [18].

9. Asymptotic Cones

9.1. **Asymptotic cones and quasi-isometries.**

Asymptotic cones are a formalization of the idea of "zooming out to infinity" on a metric space. That is, we shrink the metric to get a sense of the large-scale behavior. This is a reversal of the idea of expanding the metric at a point to get the local behavior, as in a tangent space. The idea was developed by Gromov in [13], later expanded by van den Dries and Wilkie in [19], and has ongoing importance in geometric group theory. In this thesis we give a result about the model theoretic stability of asymptotic cones of symmetric spaces and euclidean buildings.

In this section, we define asymptotic cones as ultraproducts in our logic. We also discuss maps called quasi-isometries and see a few properties of these maps. In the next section, we will see that asymptotic cones built from symmetric spaces or euclidean buildings both yield euclidean buildings.

Of course, the original definition of asymptotic cones is not phrased in terms of L-structures and ultraproducts. The definition we give below uses the framework of our logic, but is fundamentally the same. That is, if one chases out the definition of the ultraproduct, the same construction is used as in typical presentations in modern geometric group theory texts, including the use of ultrafilters.

Definition 9.1.1. Let L be a single-sorted language, and let $(M, d, \star)$ be any L-structure. Let $\mathcal{U}$ be a nonprincipal ultrafilter on $\mathbb{N}_+$. The **asymptotic cone** (with respect to $\mathcal{U}$) of $(M, d, \star)$ is the ultraproduct

$$\prod_{\mathcal{U}} (M, \frac{d}{n}, \star)$$

where $\frac{d}{n}$ denotes the metric defined by

$$\frac{d}{n}(x, y) = \frac{1}{n} d(x, y)$$

for each $n \in \mathbb{N}_+$.

We have limited the definition to a single-sorted structure to conform more to the usual terminology, but there are straightforward generalizations of this definition. One of the advantages of our framework is that it provides a straightforward way to discuss asymptotic cones with additional structures, named subsets, functions, etc. attached. It is more appropriate in our framework to think of any ultraproduct which scales down a metric in a sort as "taking the asymptotic cone" of that sort.

There are more general definitions of asymptotic cones which allow the basepoint $\star$ to vary over the factors, but we will not discuss these. We will, however, note that if the basepoint is fixed across the factors, then it does not matter which basepoint is selected.

Proposition 9.1.2. *With L, M, d, and $\mathcal{U}$ as above, for any $\star_1, \star_2 \in M$, the asymptotic cones $\prod_{\mathcal{U}}(M, \frac{d}{n}, \star_1)$ and $\prod_{\mathcal{U}}(M, \frac{d}{n}, \star_2)$ are isomorphic.*

Proof. For convenience, write

$$C_1 = \prod_{\mathcal{U}}(M, \frac{d}{n}, \star_1)$$

and

$$C_2 = \prod_{\mathcal{U}}(M, \frac{d}{n}, \star_2).$$

The isomorphism $C_1 \to C_2$ can be defined by, for any $x \in C_1$, choosing a representative $(x_n : n \in \mathbb{N}_+)$ and sending x to $(x_n : n \in \mathbb{N}_+)/\sim$ in C_2. We have $\star_1^{C_1} \mapsto \star_2^{C_2}$ because $d(\star_1, \star_2)$ is finite and hence the ultralimit of $\frac{d}{n}(\star_1, \star_2)$ is 0. In other words, $(\star_1 : n \in \mathbb{N}_+)/\sim$ and $(\star_2 : n \in \mathbb{N}_+)/\sim$ are the same class in C_2. Bijectivity and the well-definedness of this map are checked by noting that since $d(\star_1, \star_2)$ is finite, for any sequence $(x_n : n \in \mathbb{N}_+)$ the ultralimit of $\frac{d}{n}(\star_1, x_n)$ is finite if and only if that of $\frac{d}{n}(\star_2, x_n)$ is.

Hence, both ultraproducts have the same underlying set and metric. If there are any functions or relations on the sort, the relevant properties of the isomorphism follow trivially from the map's definition. □

Here are a few basic examples of asymptotic cones.

Example 9.1.3. If $(M, d, \star)$ is bounded, then all of its asymptotic cone are isomorphic to a single point space. This is because the sequence $\frac{d}{n}(x, y)$ will have ultralimit 0 for any pair of points $x, y \in M$, hence all points are identified in the ultraproduct.

Example 9.1.4. Consider $(\mathbb{Z}, d, 0)$ where d is the usual metric $d(x, y) = |x-y|$. Then any asymptotic cone is isomorphic to $(\mathbb{R}, d, 0)$, the real line with its usual metric. An isomorphism can be given by, for any $x \in \prod_{\mathcal{U}}(\mathbb{Z}, \frac{d}{n}, 0)$, choosing a representative $(x_n : n \in \mathbb{N}_+)$, and sending x to $\lim_{\mathcal{U}} \frac{1}{n} x_n$.

Example 9.1.5. Similarly to the last example, the asymptotic cones of $(\mathbb{Z}^2, d, 0)$ are isomorphic to $(\mathbb{R}^2, d, 0)$ where both d are taken to be the usual euclidean metric. This statement is also true for some other common metrics, for example if both d are taken to be the taxi-cab metric

$$((x_1, y_1), (x_2, y_2)) \mapsto |x_1 - x_2| + |y_1 - y_2|.$$

Example 9.1.6. For any euclidean space $(\mathbb{R}^n, d, 0)$, the asymptotic cones are all isomorphic to $(\mathbb{R}^n, d, 0)$. The isomorphism is similar to the above example, sending x to $\lim_{\mathcal{U}} \frac{1}{n} x_n$.

Example 9.1.7. Let $(M, d, \star)$ be the subspace $[-1, 1] \times \mathbb{R}$ of the euclidean plane $\mathbb{R}^2$ with $\star = (0, 0)$. Its asymptotic cones are isomorphic to $\{0\} \times \mathbb{R}$ with the euclidean metric.

These last few examples suggest things like that the "large-scale geometry" of $\mathbb{Z}^n$ and $\mathbb{R}^n$ are similar, and "from a large-scale perspective" bounded spaces might as well be points. Quasi-isometries are another notion that help discuss this idea.

Definition 9.1.8. Let (M_1, d_1) and (M_2, d_2) be metric spaces, and let

$$A \in [1, \infty),$$
$$B \in [0, \infty),$$
$$C \in [0, \infty).$$

We call $f : M_1 \to M_2$ an **(A, B, C)-quasi-isometry** if it has the following two properties:

- f satisfies the following inequality for all $x, y \in M_1$
$$A^{-1} d_1(x, y) - B \leq d_2(f(x), f(y)) \leq A d_1(x, y) + B.$$
- $f(M_1)$ is C-dense in M_2. That is, for every $y \in M_2$, there is $x \in M_1$ such that $d(f(x), y) \leq C$.

Quasi-isometries can be thought of as a modification of isometries which allows for some stretching, gaps, and discontinuity, within limits. For example, $\mathbb{Z}^2$ and $\mathbb{R}^2$ with the euclidean metrics are $(1, 0, 1)$-quasi-isometric, as witnessed by the identity embedding. The map $\mathbb{R}^2 \to \mathbb{Z}^2$ defined by "rounding" points to their nearest integer lattice point (say, downward in each coordinate) is a $(1, \sqrt{2}, 0)$-quasi-isometry.

A particularly nice special case of quasi-isometries is when B and C in the definition are 0. Such an f is a uniformly continuous, bijective homeomorphism.

Definition 9.1.9. Let (M_1, d_1) and (M_2, d_2) be metric spaces, and let $A \in [1, \infty)$. We say $f : M_1 \to M_2$ is an **(A)-bi-Lipschitz homeomorphism** if f is surjective and for all $x, y \in M_1$ we have

$$A^{-1} d_1(x, y) \leq d_2(f(x), f(y)) \leq A d_1(x, y).$$

We can check the following standard result that quasi-isometries become bi-Lipschitz homeomorphisms when we take asymptotic cones with respect to the same ultrafilter. In particular, notice that if $A = 1$, the (A)-bi-Lipschitz homeomorphism in the next proposition is simply an isometry.

Proposition 9.1.10. *Let $f : M_1 \to M_2$ be an (A, B, C)-quasi-isometry between L-structures $(M_1, d_1, \star_1)$ and $(M_2, d_2, \star_2)$, and let $\mathcal{U}$ be a nonprincipal ultrafilter on $\mathbb{N}_+$. Then the asymptotic cones with respect to $\mathcal{U}$ of these two structures are (A)-bi-Lipschitz homeomorphic.*

Proof. Consider the following inequalities obtained by scaling the inequalities in the definition of being a quasi-isometry

$$A^{-1}\frac{1}{n}d_1(x, y) - \frac{1}{n}B \leq \frac{1}{n}d_2(f(x), f(y)) \leq A\frac{1}{n}d_1(x, y) + \frac{1}{n}B$$

and

$$\frac{1}{n}d_2(f(x), y) \leq \frac{1}{n}C.$$

We can use f to define a map $\hat{f}$ between the asymptotic cones by sending $x = (x_n : n \in \mathbb{N}_+)/\sim$ to $(f(x_n : n \in \mathbb{N}_+))/\sim$. This is possible because f is bounded in magnitude by $AR + B$ on the d_1-ball of radius R, and if $\lim_{\mathcal{U}} \frac{1}{n}d(x_n, y_n) = 0$, then $\lim_{\mathcal{U}} \left(A\frac{1}{n}d_1(x, y) + \frac{1}{n}B\right) = 0$, so f is well-defined.

By taking ultralimits in the above inequalities, we get the (A)-bi-Lipschitz condition for $\hat{f}$ and that $\hat{f}$ is onto. ☐

In each of the previous examples, there was only one asymptotic cone up to isomorphism. This is not generally the case. The following is one of the simplest counterexamples.

Example 9.1.11. Let $(M, d, \star)$ be the subspace $\{0\} \cup \{2^n : n \in \mathbb{N}\}$ of euclidean $\mathbb{R}$ with basepoint 0. The gist of the following is that, since the gaps between the points grow exponentially, but the asymptotic cone scales the metric down linearly, we can select ultrafilters to decide whether or not we want a point to appear with a given magnitude in the asymptotic cone.

Let $\mathcal{U}$ be an ultrafilter on $\mathbb{N}_+$ containing the set $\{2^n : n \in \mathbb{N}_+\}$. For each $k \in \{2^n : n \in \mathbb{N}_+\}$, there is a point $x \in M$ with $\frac{d}{k}(0, x) = 1$, namely we can take $x = k = 2^n$. Thus the asymptotic cone (with respect to $\mathcal{U}$) has a point x with $d(0, x) = 1$.

Let $\mathcal{V}$ be an ultrafilter on $\mathbb{N}_+$ containing the set

$$P = \{2^{n+1} + 2^n : n \in \mathbb{N}_+\}.$$

If the asymptotic cone (with respect to $\mathcal{V}$) contained a point x with $d(0, x) = 1$, then the cone would satisfy the sentence $\inf]_1^2 (d(x, 0) \approx 1)$. By the fundamental theorem of ultraproducts, this means that for some $F \in \mathcal{U}$, we have for all $k \in F$ that

$$\left(\inf_x]_1^2 (d(x, 0) \approx 1)\right)^{(M, \frac{d}{k}, \star)} \leq \frac{1}{4}.$$

The rest of this discussion involves untangling what this sentence means about points in (M, d) if we select any large k, and deriving a contradiction. Since $\mathcal{V}$ is an ultrafilter and hence $F \cap P \neq \emptyset$, the above inequality holds for arbitrarily large k of the form $2^{n+1} + 2^n$. So there is some $n \geq 1$ such that

$$\left(\inf_x]_1^2 \left(d(x,0) \approx 1\right)\right)^{(M, \frac{1}{2^{n+1}+2^n} d, \star)} \leq \frac{1}{4}.$$

Evaluating this and adjusting for the scaling on the metric so that we can refer to distances in the base space $(M, d, \star)$ yields

$$\int_{2^{n+1}+2^n}^{2^{n+2}+2^{n+1}} \inf_{x \in B_\rho(M,d)} \left(\frac{||x||}{2^{n+1}+2^n} \approx 1\right) d\rho \leq \frac{1}{4}.$$

This implies that there is a point $x \in B_{2^{n+2}+2^{n+1}}(M, d)$ with $\left|\frac{||x||}{2^{n+1}+2^n} - 1\right| \leq \frac{1}{4}$, but this last inequality implies $2^{n+1} < ||x|| < 2^{n+2}$ which is impossible due to the definition of M. So this asymptotic cone must not contain a point x such that $d(x, 0) = 1$.

We have shown that the asymptotic cones with respect to $\mathcal{U}$ and $\mathcal{V}$ are not isometric, hence not isomorphic.

9.2. **Spaces that arise as asymptotic cones.** The previous example suggests a way to construct nearly arbitrary spaces as asymptotic cones, and moreover to impose heavy dependence on the ultrafilter. A brief summary of the last example is to construct the space as a union of

- repeated (and scaled up) copies of what one wants to appear under $\mathcal{U}$,
- interspersed with repeated (and scaled up) copies of what one wants to appear under $\mathcal{V}$.

In the above example, what we wanted to appear in one case was a single point, and what we wanted in the other case was nothing. We can generalize this to more than just a scattering of points. A consequence of the next result is that for any structure $(M, d, \star)$, the theory of $(M, d, \star)$ appears as the theory of some asymptotic cone.

An intuitively straightforward but technically annoying to discuss generalization of the following argument could also show that for any countable collection of structures $(M_n, d_n, \star_n)$, there is a fixed space N such that for each n, there is an ultrafilter $\mathcal{V}_n$ such that the $\mathcal{V}_n$ asymptotic cone of N is isometric to $M_n^{\mathcal{U}}$. For this, one just has to "interleave" the annuli in the following construction, for example by always having the p_n^n-th annulus correspond to M_n, where p_n is the n-th prime.

Proposition 9.2.1. *Let L be a single-sorted language with no additional symbols, let $(M, d_M, \star_M)$ be an L-structure, and let $\mathcal{U}$ be a nonprincipal ultrafilter on $\mathbb{N}_+$.*

Then there exists an L-structure $(N, d_N, \star_N)$ and a nonprincipal ultrafilter $\mathcal{V}$ on $\mathbb{N}_+$ such that the ultrapower $M^{\mathcal{U}}$ is isometric to the asymptotic cone $\prod_{\mathcal{V}}(N, \frac{1}{n} d_N, \star_N)$.

Proof. First we illustrate an idea in the construction that follows. Given any annulus $A = B_{r_2}(M) - B_{r_1}(M)$, which has inner radius r_1 and outer radius r_2, we can think of it as a subspace which has empty $B_{r_1}(A)$ and is entirely contained in $B_{r_2}(A)$. if we multiply the metric on this subspace by k, we obtain a space A' which has an empty $B_{k\cdot r_1}(A')$ and is entirely contained in $B_{k\cdot r_2}(A')$.

The rest of this proof consists of repeating the above idea to build N as a union of copies of annular subspaces of M, so that for arbitrarily large n, the space $(N, \frac{1}{n}d, \star)$ "looks like" $B_n(M) - B_{\frac{1}{n}}(M)$. Then we just choose an ultrafilter containing a sequence of such arbitrarily large n.

Let N_1 be just $\{\star\}$, and define $k_1 = 1$. Let N_2 be $B_2(M) - B_{\frac{1}{2}}(M)$ and define $k_2 = 2$. We rescale the metric on N_2 by multiply by k_2. That is, the metric on N_2 is $k_2 \cdot d(x, y)$ where we view x, y as elements of M. Notice that N_2 has empty $B_{k_2\cdot\frac{1}{2}}(N) = B_{k_1\cdot 1}(N)$ and is contained within $B_{k_2\cdot 2}(N)$. In general, define N_n to be $B_n(M) - B_{\frac{1}{n}}(M)$, and define k_n so that $k_n \cdot \frac{1}{n} = k_{n-1}\cdot(n-1)$. We scale the metric on N_n by multiplying by k_n. Observe that N_n has empty $B_{2^n\cdot\frac{1}{n}}(N)$ and is contained within $B_{2^n\cdot n}(N)$. This explains the selection of k_n. At each step, we are ensuring that the annulus has an inner radius at least as large as the outer radius of the previous annulus. We sumarize the first few steps below in a table.

n	k_n	N_n has empty radius...	N_n is contained in radius...
1	1	$-$	$-$
2	2	1	4
3	12	4	36
4	144	36	576
n	k_n	$k_n \cdot \frac{1}{n}$	$k_n \cdot n$
$n+1$	k_{n+1}	$k_{n+1}\cdot\frac{1}{n+1} = k_n\cdot n$	$k_{n+1}\cdot n$

Now, to construct N, we take the disjoint union $\bigsqcup N_n$. To define the metric on N, use the metrics above for pairs of points within a common N_n, and for pairs in different annuli, say $x \in N_n$ and $y \in N_m$, define

$$d_N(x, y) = k_n d(x, \star) + k_m d(y, \star)$$

where we view x, y as points in M to apply d.

Consider $(N, \frac{1}{k_n}d_N, \star)$ for any of the k_n appearing above. Since

$$B_{k_n\cdot n}(N, d_N) - B_{k_n\cdot\frac{1}{n}}(N, d_N)$$

is an up-scaled copy of $B_n(M) - B_{\frac{1}{n}}(M)$, the subspace

$$B_n(N, \frac{1}{k_n}d_N) - B_{\frac{1}{n}}(N, \frac{1}{k_n}d_N)$$

is isometric to $B_n(M) - B_{\frac{1}{n}}(M)$.

This suggests an isometry f between $M^{\mathcal{U}}$ and the ultraproduct $\prod_{\mathcal{U}}(N, \frac{1}{k_n}d_N, \star)$. Any sequence $(x_n : n \in \mathbb{N}_+)$ representing a class in $M^{\mathcal{U}}$ should be sent to

$(x'_n : n \in \mathbb{N})/\sim$ in this ultraproduct, where $x'_n \in N_n \subseteq N$ is the point corresponding to x_n if $x_n \in B_n(M) - B_{\frac{1}{n}}(M)$, and is $\star$ otherwise.

We will check first that if we have two elements of $M^{\mathcal{U}}$,

$$x = (x_n : n \in \mathbb{N}_+)/\sim$$
$$y = (y_n : n \in \mathbb{N}_+)/\sim,$$

then their distance is preserved by this map. Since $\lim_{\mathcal{U}} d(x_n, \star) = ||x||$ and $\lim_{\mathcal{U}} d(y_n, \star) = ||y||$ are both finite, either their representative sequences satisfy $x_n = x'_n$ and $y_n = y'_n$ for all n in some $F \in \mathcal{U}$, or at least one of the representative sequences satisfies $x_n = \star$ or $y_n = \star$ for all n in some $F \in \mathcal{U}$. In the first case, the distance is preserved because the isometry between the annuli ensures $d(x_n, y_n) = \frac{1}{k_n} d_N(x'_n, y'_n)$ for all n in this $F \in \mathcal{U}$. In the second case, the distance is preserved because either both x, y are mapped to the basepoint of the ultraproduct, or else just one is (say y) while the other satisfies $x_n = x'_n$ for all $n \in F \in \mathcal{U}$. In that remaining case, we use that the construction of N shows $\frac{1}{k_n} d_N(x'_n, \star) = d(x_n, \star)$.

Similar considerations show the well-definedness of f. To see that f is surjective, note that if $z = (z_n : n \in \mathbb{N}_+)/\sim$ is an element of the ultraproduct $\prod_{\mathcal{U}}(N, \frac{1}{k_n} d_N, \star)$ other than the basepoint (which is a trivial case), then we have

$$||z|| \;=\; \lim_{\mathcal{U}} \frac{1}{k_n} d_N(z_n, \star) > 0.$$

So there is some $F \in \mathcal{U}$ such that for all $n \in F$ we have $\frac{1}{k_n} d_N(z_n, \star)$ between $\frac{1}{n}$ and n. Let $x \in M^{\mathcal{U}}$ be the class of the sequence $(\hat{z_n} : n \in \mathbb{N})$, where $\hat{z_n}$ denotes the element in M corresponding to the point $z_n \in N$. The map f sends x to $x' = (\hat{z_n}' : n \in \mathbb{N})/\sim$ in the ultraproduct. But we have just shown that $\hat{z_n}' = z_n$ for $n \in F$. So $x' = z$ in $\prod_{\mathcal{U}}(N, \frac{1}{k_n} d_N, \star)$.

To complete the proof, we just need to note that $\prod_{\mathcal{U}}(N, \frac{1}{k_n} d_N, \star)$ can be realized as an asymptotic cone $\prod_{\mathcal{V}}(N, \frac{1}{n} d_N, \star)$ by using an ultrafilter $\mathcal{V}$ defined to include $\{k_n : n \in F\}$ for each $F \in \mathcal{U}$. $\square$

9.3. **An intermediate value theorem.** In this section we prove that for many natural spaces M, if M has two asymptotic cones which have distinct values $r_1 < r_2$ for some sentence, then we can find cones that obtain every intermediate value for that sentence. The assumption we need on M is that each ball $\bar{B}_n(M)$ is sufficiently dense in $\bar{B}_{n+1}(M)$. This will be satisfied in any geodesic space, for example.

First we prove a lemma which says the values of formulas have to change very slowly when interpreted in sufficiently scaled down copies of M.

Lemma 9.3.1. *Let L be the language with no additional functions or relations, and let $(M, d, \star)$ be an L-structure such that for some $D \geq 0$, for all $n \in \mathbb{N}_+$, the ball $\bar{B}_n(M, d)$ is D-dense in $\bar{B}_{n+1}(M, d)$. That is, every point in $\bar{B}_{n+1}(M, d)$ is at most distance D from a point in $\bar{B}_n(M, d)$.*

For every $r, \varepsilon > 0$ and formula $\phi(x)$, there is $N \in \mathbb{N}$ such that whenever $n \geq N$ and a is a tuple in $B_{r \cdot n}(M, d)$, we have

$$\left|\phi^{(M, \frac{d}{n}, \star)}(a) - \phi^{(M, \frac{d}{n+1}, \star)}(a)\right| \leq \varepsilon.$$

Proof. We prove this by induction on formulas. The atomic formulas are just of the form $d(p, p')$ with $p, p' \in B_{r \cdot n}(M, d)$. Here we have

$$\begin{aligned} \left|\frac{d(p,p')}{n} - \frac{d(p,p')}{n+1}\right| &= \frac{d(p,p')}{n(n+1)} \\ &\leq \frac{2rn}{n(n+1)} \\ &= \frac{2r}{n+1} \end{aligned}$$

which is $\leq \varepsilon$ for sufficiently large n.

Suppose $\psi(x)$ is of the form $u(\phi_1(x), \ldots, \phi_k(x))$ where each ϕ_i has already been checked. Whenever $x \in B_{r \cdot n}(M, d)$, we have $x \in B_r(M, \frac{d}{n})$. Each ϕ_i is controlled since it is a formula, so we obtain a uniform bound b such that

$$|\phi_i^{(M, \frac{d}{n}, \star)}(x)| \leq b$$

and

$$|\phi_i^{(M, \frac{d}{n+1}, \star)}(x)| \leq b$$

for all $x \in B_{r \cdot n}(M, d)$. By uniform continuity of u on $[0, b]^k$, there is some δ such that inputs within δ of each other are mapped by u to within ε of each other. By hypothesis, there are N_i after which each corresponding ϕ_i have interpretations in $(M, \frac{d}{n}, \star)$ and $(M, \frac{d}{n+1}, \star)$ differing by at most δ. So the claim for ψ holds using

$$N = \max\{N_i : i = 1, \ldots, k\}.$$

Suppose $\psi(x)$ is of the form $\sup_y]_{r'}^{r} \phi(x, y)$ with $r' < r$ and where $\phi(p, p')$ has already been checked for each pair p, p' in $B_{r \cdot n}(M, d)$. By approximating $\sup_y]_{r'}^{r} \phi(x, y)$ with Riemann sums as in Lemma 2.7.6, there are finitely many $\rho \in [r', r]$ such that it suffices to know that $\sup_{y \in B_\rho(M, \frac{d}{n})} \phi(p, y)$ and $\sup_{y \in B_\rho(M, \frac{d}{n+1})} \phi(p, y)$ can be made as close as desired for sufficiently large n. But for any given $\rho \in [r', r]$, this follows from the inductive hypothesis and the D-density assumption. □

Proposition 9.3.2. *Let L be the language with no additional functions or relations, and let $(M, d, \star)$ be an L-structure such that for some $D \geq 0$, for all $n \in \mathbb{N}_+$, the ball $\bar{B}_n(M, d)$ is D-dense in $\bar{B}_{n+1}(M, d)$.*

Let C_1 and C_2 be two asymptotic cones of M. Let σ be an L-sentence, let $r_1 < r_2$ be reals, and suppose that

$$\sigma^{C_1} = r_1,$$
$$\sigma^{C_2} = r_2.$$

Then for any r with $r_1 < r < r_2$, there is an asymptotic cone C of M where

$$\sigma^C = r.$$

Proof. By applying Łoś's Theorem to both of C_1 and C_2, σ must be arbitrarily close to each r_1 and r_2 infinitely often. But by the previous proposition, for any $\varepsilon > 0$, we have that eventually $\sigma^{(M, \frac{d}{n+1}, \star)}$ is within ε of $\sigma^{(M, \frac{d}{n}, \star)}$. So $\sigma^{(M, \frac{d}{n}, \star)}$ must be arbitrarily close to r infinitely often.

This means there is a sequence $(a_k)_{k \in \mathbb{N}}$ such that

$$\left| \sigma^{(M, \frac{d}{a_k}, \star)} - r \right| \leq \frac{1}{k}.$$

By choosing an ultrafilter $\mathcal{U}$ containing $\{a_k : k \in \mathbb{N}\}$ and applying Łoś's Theorem to the asymptotic cone of M with respect to $\mathcal{U}$, we obtain the result. □

Corollary 9.3.3. *Let M be as in the previous proposition. Then M either has 1 or $\geq 2^{\aleph_0}$ asymptotic cones up to elementary equivalence.*

Proof. If there are two cones which are not elementary equivalent, then there is a sentence σ which has a value r_1 in one cone and r_2 in the other cone. By the above proposition, we obtain a cone with a distinct theory for each $r \in [r_1, r_2]$. □

9.4. **Ultraproducts of symmetric spaces and euclidean buildings.** In this section, we discuss the theorems from [14] showing that, provided we restrict ourselves to $(\mathbb{E}, W)$ spaces,

- the asymptotic cones of symmetric spaces of noncompact type are euclidean buildings, and
- arbitrary ultraproducts of euclidean buildings are euclidean buildings.

In both cases, the atlas on the asymptotic cone is induced in a straightforward way from the $(\mathbb{E}, W)$ space atlas on the base space or spaces. The reason we chose to use the $(\mathbb{E}, W)$ atlas approach based on Kleiner and Leeb's presentation of euclidean buildings is that, as they mention, this makes the euclidean building definition very easy to check in the asymptotic cones.

Earlier, we mentioned the intuition suggested by Kleiner and Leeb that, while symmetric spaces fail to have the angle rigidity property that euclidean buildings have, symmetric spaces do have this angle rigidity "at infinity". This is a reference to the nice asymptotic behavior of geodesics in symmetric spaces, which was needed obtain the spherical building at infinity. The intuition for the symmetric space case is that since the metric is scaled down in the asymptotic cone construction, this asymptotic behavior of rays is pulled to the origin, so that the nice asymptotic behavior of pairs of rays in M becomes a nice initial behavior between the corresponding rays in the asymptotic cone of M.

It may be helpful to recall that asymptotic cones of $\mathbb{R}^k$ are all isometric to $\mathbb{R}^k$. When we talk about the ray γ' in $(M, \frac{d}{n}, \star)$ corresponding to a ray γ in $(M, d, \star)$, we mean the same subspace of M, but the map $\gamma : I \to M$ must be reparametrized to get an isometry $\gamma' : I \to M$. The rays that appear in an asymptotic cone over symmetric spaces of noncompact type or over euclidean buildings will arise as ultraproducts of rays (viewed as subspaces) in the base space. This is a consequence of the definability of rays in the class of CAT(0) spaces.

Because $\mathbb{R}^k$ stays essentially the same under ultrapowers and asymptotic cones, and $\Delta_{(\mathbb{E},W)}$ is compact and stays the same under ultrapowers, the proofs are almost immediate and just require noting that the natural choices satisfy the expected properties. For example, the ultralimits of charts yield charts, and $d_\triangle$ distances are bounded by initial angles because the inequality is preserved under ultralimits.

In both cases, the interesting and substantial property to check is the angle rigidity property, much like when we wanted to see the spherical buildings at infinity. We are able to take arbitrary ultraproducts of euclidean buildings in this result because each already satisfies the property. For symmetric spaces, we have to take the asymptotic cone to get the property.

Proposition 9.4.1. *For $n \in \mathbb{N}$, let M_n be an $(\mathbb{E}, W)$ euclidean building. Then every ultraproduct $\prod_{\mathcal{U}} M_n$ is also an $(\mathbb{E}, W)$ euclidean building.*

Proof. This is Theorem 5.1.1 of [14], translated to our notation. □

Proposition 9.4.2. *Let $(M, d, \star)$ be a symmetric space of noncompact type with an $(\mathbb{E}, W)$ atlas. Then every asymptotic cone $\prod_{\mathcal{U}}(M_n, \frac{1}{n}d, \star)$ is an $(\mathbb{E}, W)$ euclidean building.*

Proof. Similarly, this is Theorem 5.2.1 of [14]. □

10. Definability of Building Structure

To summarize our current position, consider the following.

- Each symmetric space of noncompact type $(M, d, \star)$ has an associated maximal $(\mathbb{E}, W)$ atlas determined by the geometry of flats in (M, d).
- This maximal $(\mathbb{E}, W)$ atlas induces a euclidean building structure on every asymptotic cone of $(M, d, \star)$.
- All of these spaces have an associated spherical building at infinity, which we view as a structure on $\Gamma_\star$, the rays at the origin.
- We have already seen that $\Gamma_\star$ is definable.
- In these spherical buildings, we can construct a projective plane and a field by working with the incidence structure on vertices.

In the upcoming sections, we show how continuous logic captures this incidence structure of the spherical building, and lets us construct the projective plane and field just mentioned. We work across the class of asymptotic cones of a given symmetric space of noncompact type $(M, d, \star)$. Our language will only use the metric symbol, basepoint, and a finite number of constants used to label some vertices.

Recall that since we are dealing with CAT(0) spaces with extension of segments to rays, we already have definability of $\Gamma_\star N$ in the class of spaces we are working in. So we can quantify over $\Gamma_\star N$. Moreover, since $\Gamma_\star N$ is bounded, this quantification is exact. Our focus is on subsets of $\Gamma_\star N$ which have special meaning in the spherical building at infinity.

10.1. **Definability of orbits.** The $(\mathbb{E}, W)$ atlas in a euclidean building N provides a map θ that takes any $\gamma \in \Gamma_\star N$ and assigns an element of the anisotropy polyhedron $\Delta_{(\mathbb{E},W)}$. In this section, we show the definability of the set of rays sharing a given image under θ. In particular, this means we get definability of the set of vertices of a given label (as in section 8.9) in the spherical building at infinity.

The next definition introduces notation for preimages under θ in order to make our discussion easier.

Definition 10.1.1. Let N be an $(\mathbb{E}, W)$ space, so that we have a map θ and an anisotropy polyhedron. Let $\Theta \in \Delta_{(\mathbb{E},W)}$. We will write $\Gamma^\Theta_\star(N)$ to mean the set

$$\{\gamma \in \Gamma_\star(N) : \theta(\gamma) = \Theta\}$$

of rays starting at $\star$ and with $\Delta_{(\mathbb{E},W)}$-direction Θ.

Recall that we fix an ordered set of labels $1, \dots, k$ on the vertices of $\Delta_{(\mathbb{E},W)}$ when $\mathbb{E} = \mathbb{R}^k$, as in 8.9, and that this labeling induces a labeling on the vertices of the spherical building by taking preimages under θ.

In this and the next few sections, we will be talking about symmetric spaces with a selection of constant rays. To help state the results more succinctly, we make the following definition.

Definition 10.1.2. By a **labeled symmetric space**, we will mean a symmetric space of noncompact type $(M, d, \star)$ with rank $K \geq 2$, equipped with its maximal $(\mathbb{E}, W)$ atlas, and with $\Theta_1, \dots, \Theta_K$ the ordered vertices of $\Delta_{(\mathbb{E},W)}$. We associate to each Θ_k a ray $\gamma_k \in \Gamma^{\Theta_k}_\star(M)$, which we will think of as an $\mathbb{N}$-indexed sequence.

Thus, a labeled symmetric space is a metric structure $(M, d, \star, \gamma_1, \dots, \gamma_K)$. We will be looking at structures arising from asymptotic cones of these, and extensions of them by a few more constants. In every case, we will work in an appropriate language L with at least K many constant symbols for the rays $\gamma_1, \dots, \gamma_K$, and possibly more constants for rays, so that our labeled symmetric space is an L-structure. The controllers, in particular the bounds for λ, can be found by noting that an $\mathbb{N}$-indexed ray in $\Gamma_\star$ is contained in $\prod_{n\in\mathbb{N}} \bar{B}_n(M)$.

Finally, recall that asymptotic cones of these spaces will be euclidean buildings. We can extend the asymptotic cone construction to the labeled symmetric space structure, but this requires a note about the rays. Namely, if we just took

$$\prod_{\mathcal{U}} (M, \frac{d}{n}, \star, \gamma_1, \dots, \gamma_K)$$

then the rays γ_k would degenerate to $(\star : i \in I)/\sim$. But, there is a clear correspondence between rays in (M, d) and rays in $(M, \frac{d}{n})$ given by rescaling the rays.

Definition 10.1.3. Given a ray γ in (M, d), we define the n**-th rescaling** of γ to be the ray $n\gamma$ in $(M, \frac{d}{n})$ defined by

$$n\gamma(r) = \gamma(rn)$$

for all $r \in \mathbb{R}_{\geq 0}$.

Now we can extend the definition of asymptotic cones.

Definition 10.1.4. If M is a labeled symmetric space $(M, d, \star, \gamma_1, \ldots, \gamma_K)$ and $\mathcal{U}$ is a nonprincipal ultrafilter over $\mathbb{N}$, then the asymptotic cone of M with respect to $\mathcal{U}$ is the structure

$$\prod_{\mathcal{U}} (M, \frac{d}{n}, \star, n\gamma_1, \ldots, n\gamma_K).$$

Of course, this asymptotic cone is then a "labeled euclidean building", in the sense that the resulting rays still satisfy

$$\gamma_k \in \Gamma_\star^{\Theta_k}.$$

Again, note that we work in a language L having nothing but constant symbols for the rays $\gamma_1, \ldots, \gamma_K$.

Theorem 10.1.5. *Let M be a labeled symmetric space, and let $\Theta \in \Delta_{(\mathbb{E},W)}$.*
Let $\mathcal{C}$ be the class of asymptotic cones $(N, d, \star, \gamma_1, \ldots, \gamma_K)$ of M.
Then $\Gamma_\star^\Theta$ is L-definable in the class $\mathcal{C}$. That is, there are L-formulas depending only on $\mathcal{C}$ which define $\Gamma_\star^\Theta(N)$ for each $N \in \mathcal{C}$.

Proof. Since we are checking definability of a set of sequences, we need to show that every finite projection $\pi_I(\Gamma_\star^\Theta)$ is definable in the class.

Pick any N in $\mathcal{C}$. Let $I \subseteq \mathbb{N}$ be finite, and write P for the projection $\pi_I(\Gamma_\star^\Theta)$ of $\Gamma_\star^\Theta$ onto its coordinates in I. That is, P consists of the I-tuples $(a_i : i \in I)$ which are projections of rays $\gamma \in \Gamma_\star(N)$ satisfying $\theta(\gamma) = \Theta$.

First we find a formula $\phi(x_i : i \in I)$ with P as its zero set. To do this, we use our knowledge of the finite sets $D(\Theta, \Theta_k)$. The key insights are the following.

- $\theta(\gamma) = \Theta$ iff for each $k \in \{1, \ldots, K\}$, the angle between γ and γ_k is one of the finitely many angles in $D(\Theta, \Theta_k)$. That is, we can triangulate the position of $\theta(\gamma)$ in the polytope $\Delta_{(\mathbb{E},W)}$ by considering the angles $\theta(\gamma)$ makes with the vertices $\Theta_1, \ldots, \Theta_K$.
- We can test if the angle between γ and γ_k is in $D(\Theta, \Theta_k)$ by finding subsegments $\gamma([t_0, t_1])$ and $\gamma_k([t_0, t_1])$ which occopy a common apartment and diverge according to an angle in $D(\Theta, \Theta_k)$.

- We can find such subsegments using finitely many fixed t by exploiting what we know about $d(\gamma(t), \gamma_k(t))$ from the piecewise linearity Proposition 8.6.4 and its corollary.

To help standardize our approach, let $J \subseteq \mathbb{N}$ be the finite set $I \cup \{0, 1, \ldots, 2n+1\}$ where n is the number of reflections in $W_\star$. Note that n is an upper bound for $D(\theta(\gamma), \Theta_k)$ over all γ and $k \in \{1, \ldots, K\}$. Consequently, each function $d(\gamma(t), \gamma_k(t))$ consists of at most n many linear pieces. We will sample the rays at $t = 0, 1, \ldots, 2n+1$ and use a pigeon-hole argument to estimate the slope of at least one of these linear pieces. Recall that we can quantify exactly over $\pi_J(\Gamma_\star N)$.

Let ϕ be given by

$$\inf_{(y_j : j \in J) \in \pi_J(\Gamma_\star N)} \left(\sum_{i \in I} d(x_i, y_i) + \sum_{k \in \{1,\ldots,K\}} \prod_{j \in \{0,\ldots,n\}} (\psi_{k,j} + \psi'_{k,j}) \right)$$

where $\psi_{k,j}$ and $\psi'_{k,j}$ are the formulas defined below. Because of the quantification, the variables y_j range over points on rays $\gamma \in \Gamma_\star N$. For convenience below, we will refer to such rays γ, though we will only be discussing the points $y_0, \ldots, y_{2n+1}$. Recall that each γ_k is a constant sequence and we are given its coordinates, which we write below in the form $\gamma_k(j)$.

The $\sum_{i \in I} d(x_i, y_i)$ term just measures the distance of a tuple (x_i) to the relevant projection (y_i) of a ray γ.

The formula $\psi_{k,j}$ measures the flatness of the segments $\gamma([2j, 2j+2])$ and $\gamma_k([2j, 2j+2])$ by comparing two secant lines of the function $d(\gamma(t), \gamma_k(t))$ on $[2j, 2j+2]$. It compares the slope of the secant of the first half to that of the whole length.

$\psi_{k,j}$ is defined to be

$$\frac{|d(y_{2j}, y_{2j+1}) - d(\gamma_k(2j), \gamma_k(2j+1))|}{1} \approx \frac{|d(y_{2j}, y_{2j+2}) - d(\gamma_k(2j), \gamma_k(2j+2))|}{2}$$

Notice that $\psi_{k,j}$ is 0 exactly when $d(\gamma(t), \gamma_k(t))$ is linear, since we already know this function is convex.

The formula $\psi'_{k,j}$ measures the divergence of γ and γ_k by comparing the slope of the secant of the function $d(\gamma(t), \gamma_k(t))$ to the possible slopes corresponding to the angles $\alpha \in D(\Theta, \Theta_k)$.

$\psi'_{k,j}$ is defined to be

$$\prod_{\alpha \in D(\Theta, \Theta_k)} \left(\frac{|d(y_{2j}, y_{2j+1}) - d(\gamma_k(2j), \gamma_k(2j+1))|}{1} \approx 2\sin(\alpha/2) \right)$$

Notice that $\psi'_{k,j}$ is 0 exactly when this slope of the secant over $[2j, 2j+1]$ corresponds to one of the $\alpha \in D(\Theta, \Theta_k)$. When we additionally have $\psi_{k,j} = 0$, this slope of the secant is actually the slope of a linear piece of $d(\gamma(t), \gamma_k(t))$.

We will now verify that P is the zero set of ϕ.

First we show that P is a subset of the zero set of ϕ. Suppose $(x_i : i \in I) \in P$. That is, (x_i) is the I-projection of some ray γ with $\theta(\gamma) = \Theta$. Let $(y_j : j \in J)$ be the J-projection of this same γ. We will check that the inf in ϕ is 0 and this is witnessed by the y we have just identified. The $\sum d$ term of ϕ is 0 because $x_i = y_i$ for $i \in I$. For the $\sum_k \prod_j$ term, we will argue that for each $k \in \{1, \ldots, K\}$ there is at least one j where $\psi_{k,j}$ and $\psi'_{k,j}$ are both 0.

Let $k \in \{1, \ldots K\}$. The function $d(\gamma(t), \gamma_k(t))$ consists of at most n many linear pieces. Our product is over $(n+1)$ many values of j and involves $(n+1)$ many intervals $[2j, 2j+2]$ with pairwise disjoint interiors. So $d(\gamma(t), \gamma_\Theta(t))$ must be linear on at least one of these intervals. For such an interval $[2j, 2j+2]$, the slope of the secant line over $[2j, 2j+1]$ is exactly the slope of the secant line over $[2j, 2j+2]$. This implies $\psi_{k,j} = 0$. Since this is the slope of a linear part of the function and we have assumed $\theta(\gamma) = \Theta$, this slope must be $2\sin(\alpha/2)$ for some $\alpha \in D(\Theta, \Theta_k)$. This implies $\psi'_{k,j} = 0$.

We have shown that P is contained in the zero set of ϕ. We now handle the reverse containment. Suppose $\phi(x_i : i \in I) = 0$. Since asymptotic cones are countable ultraproducts and hence $\aleph_1$-saturated, the inf is realized by some $(y_j) \in \pi_J(\Gamma_\star N)$. This (y_j) is the J-projection of a ray γ such that for all $k \in \{1, \ldots, K\}$, there is some j with $\psi_{k,j} = \psi'_{k,j} = 0$.

Since $\psi_{k,j} = 0$, the function $d(\gamma(t), \gamma_k(t))$ is linear on $[2j, 2j+2]$. Since $\psi'_{k,j} = 0$, the slope of $d(\gamma(t), \gamma_k(t))$ in this interval is $2\sin(\alpha/2)$ for some $\alpha \in D(\Theta, \Theta_k)$. So, $\gamma \in \Gamma^\Theta_\star$. Since the $\sum d$ term in ϕ is 0, we conclude that $(x_i : i \in I)$ is equal to $\pi_I(\gamma) \in P$.

To complete the proof we will show that we can obtain the distance to P using ϕ and the ε-δ lemma 3.9.6.

Suppose $\phi(x_i : i \in I) \leq \delta$ for some $\delta > 0$. The inf in ϕ is witnessed by some $(y_j : j \in J)$ which is the projection $\pi_J(\gamma)$ of some $\gamma \in \Gamma_\star N$ and satisfies

$$\left(\sum_{i \in I} d(x_i, y_i)\right) + \sum_{k \in \{1,\ldots,K\}} \prod_{j \in \{0,\ldots,n\}} (\psi_{k,j} + \psi'_{k,j}) \leq \delta.$$

Since both terms are nonnegative, we get

$$\sum_{i \in I} (d(x_i, y_i)) \leq \delta,$$

and we get for all $k \in \{1, \ldots, K\}$ that

$$\prod_{j \in \{0,\ldots,n\}} (\psi_{k,j} + \psi'_{k,j}) \leq \delta.$$

This last inequality implies that $(\psi_{k,j} + \psi'_{k,j}) \leq \delta^{1/(n+1)}$ for at least one j, and hence $\psi_{k,j}, \psi'_{k,j} \leq \delta^{1/(n+1)}$ for at least one $j \in \{0, \ldots, n\}$. Fix such a j for each k.

For each $k \in \{1, \ldots, K\}$, we argue as follows. Since each $d(\gamma(t), \gamma_k(t))$ is a piecewise linear, nondecreasing function of t, there must be a subinterval S of

$[2j, 2j+2]$ where $d(\gamma(t), \gamma'(t))$ has a slope between the slopes m_1 and m_2 of the secant lines.

$$\begin{aligned} m_1 &= \frac{|d(y_{2j}, y_{2j+1}) - d(\gamma_{2j}, \gamma_{2j+1})|}{1} \\ m_2 &= \frac{|d(y_{2j}, y_{2j+2}) - d(\gamma_{2j}, \gamma_{2j+2})|}{2} \end{aligned}$$

The bound on ψ_j implies that

$$|m_1 - m_2| \leq \delta^{1/(n+1)},$$

and so the slope on S is within $\delta^{1/(n+1)}$ of m_1. The bound on ψ'_j means that m_1 is within $\delta^{1/(n+1)}$ of one of the slopes $2\sin(\alpha/2)$ correspondng to an angle $\alpha \in D(\Theta, \Theta_k)$. Thus, the slope on S is within $2\delta^{1/(n+1)}$ of one of the values $2\sin(\alpha/2)$.

Since γ exists in some apartment, this shows that by selecting δ small enough, we can ensure γ is as close as we want to another ray γ' which has a segment where $d(\gamma'(t), \gamma_k(t))$ is linear and has slope $2\sin(\alpha/2)$ for some $\alpha \in D(\Theta, \Theta_k)$. That is, for sufficiently small δ, we have γ arbitrarily close to $\Gamma_\star^\Theta$. Since $(x_i : i \in I)$ is at most δ away from $\pi_I(\gamma)$, this shows we can make $d((x_i), \pi_I(\Gamma_\star^\Theta))$ arbitrarily small by selecting δ small enough. $\square$

As a corollary, we note that because any $\gamma \in \Gamma_\star$ which lies in the interior of a chamber or face determines that chamber or face uniquely, we can essentially use the above result to quantify over chambers and faces of the spherical building at infinity by quantifying over $\Gamma_\star^\Theta$ for some fixed Θ in the corresponding simplex of $\Delta_{(\mathbb{E},W)}$. We will not need to do this for our application, but we note it because it indicates that we are interpreting the simplicial spherical building structure. In the end of the next section, we will note that we can obtain the containment relation on these simplices.

For our application, our focus is on the vertices and the incidence relation given by containment in a common chamber.

10.2. **Definability of the incidence relation.** Spherical buildings have an incidence relation on their vertices. Two vertices are called incident when they occupy a common face or chamber. Equivalently, since any two vertices are contained in some apartment, we can say two distinct vertices v_1, v_2 are incident when their distance is minimal among $D(\theta(v_1), \theta(v_2))$.

When discussing the spherical building at infinity of a euclidean building, that can be written as follows. Taking Θ and Θ' to be distinct points in $\Delta_{(\mathbb{E},W)}$, two rays $\gamma \in \Gamma_\star^\Theta$ and $\gamma' \in \Gamma_\star^{\Theta'}$ are incident if there is a euclidean apartment containing both rays and the angle $\angle_\star(\gamma_1, \gamma_2)$ is minimal among $D(\Theta, \Theta')$.

Yet another equivalent characterization is that $\gamma \in \Gamma_\star^\Theta$ and $\gamma' \in \Gamma_\star^{\Theta'}$ are incident if the function $d(\gamma_1(t), \gamma_2(t))$ is just the line with slope $2\sin(\alpha/2)$ where α is the minimal value in $D(\Theta, \Theta')$. This characterizes the incidence relation purely in terms of the distances between the rays.

Theorem 10.2.1. *Let M be a labeled symmetric space, and let Θ, Θ' be a distinct pair in $\Delta_{(\mathbb{E},W)}$.*

Let $\mathcal{C}$ be the class of asymptotic cones $(N, d, \star, \gamma_1, \ldots, \gamma_K)$ of M.

Then the set of vertices $\gamma' \in \Gamma_\star^{\Theta'}$ which are incident with a given $\gamma_\Theta \in \Gamma_\star^{\Theta}$ is L-definable over γ_Θ in the class $\mathcal{C}$.

Proof. We show that the finite projections are definable. Let $I \subseteq \mathbb{N}$ be finite. Recall that we can quantify exactly over $\Gamma_\star^{\Theta}$ and $\Gamma_\star^{\Theta'}$ by definability result of the previous section and the fact that each finite projection of these sets is bounded.

Consider the formula ϕ given by

$$\inf_{(y_i : i \in I) \in \pi_I(\Gamma_\star^{\Theta'})} \left(\left(\sum_{i \in I} d(x_i, y_i) \right) + \sum_{i \in I} \left(d(y_i, \gamma_\Theta(i)) \approx i \cdot 2\sin(\alpha/2) \right) \right)$$

where α is the minimal value of $D(\Theta, \Theta')$.

First, we check that ϕ has the right zero set. Suppose $(x_i : i \in I)$ is the I-projection of a ray $\gamma' \in \Gamma_\star^{\Theta'}$ incident with γ_Θ. As discussed in the paragraph introducing this section, this means $d(\gamma'(t), \gamma_\Theta(t))$ is a linear function with slope $2\sin(\alpha/2)$. Hence (x_i) itself witnesses that the inf in ϕ is 0.

For the converse, suppose $\phi(x_i : i \in I) = 0$. Then since asymptotic cones are $\aleph_1$-saturated, there is some (y_i) realizing the inf. This (y_i) is the I-projection of a ray $\gamma \in \Gamma_\star^{\Theta'}$, and because the $\sum_{i \in I}$ term in ϕ is 0, we know the function $d(\gamma(t), \gamma_\Theta(t))$ is linear and has slope $2\sin(\alpha/2)$ on at least the interval

$$S = [0, \max(i \in I)].$$

Since the segment $\gamma(S)$ is in an apartment and can be extended to a ray which maintains this slope, we can assume without loss of generality that (y_i) and hence γ is chosen incident to γ_Θ. Since the $\sum d$ term in ϕ is 0, we have that $(x_i) = (y_i)$. Thus (x_i) is the I-projection of a ray in $\Gamma_\star^{\Theta'}$ incident with γ_Θ, as required.

Finally, we check that we can apply the ε-δ lemma 3.9.6 to complete the proof. Let $\delta > 0$ and suppose $\phi(x_i : i \in I) \leq \delta$. Witness the inf by some $(y_i) \in \pi_I(\Gamma_\star^{\Theta'})$. We have that (y_i) is the I-projection of a ray $\gamma' \in \Gamma_{\Theta'}$ and satisfies

$$\left(\sum_{i \in I} d(x_i, y_i) \right) + \sum_{i \in I} \left(d(y_i, \gamma_\Theta(i)) \approx i \cdot 2\sin(\alpha/2) \right) \leq \delta.$$

This implies

$$\sum_{i \in I} d(x_i, y_i) \leq \delta,$$

and for all $i \in I$ we also get

$$\left(d(y_i, \gamma_\Theta(i)) \approx i \cdot 2\sin(\alpha/2) \right) \leq \delta.$$

Since $\gamma' \in \Gamma_\star^{\Theta'}$, we know that the linear pieces of $d(\gamma'(t), \gamma_\Theta(t))$ must have slopes in $D(\Theta, \Theta')$. But $D(\Theta, \Theta')$ is finite, so for sufficiently small δ, the last inequality above implies that $d(y_i, \gamma_\Theta(i)) \approx i \cdot 2\sin(\alpha/2)$ is exactly 0, i.e.

$$d(y_i, \gamma_\Theta(i)) = i \cdot 2\sin(\alpha/2).$$

As we argued earlier, this means that the subsegment $\gamma'([0, \max(i \in I)])$ can be extended to a ray adjacent to γ_Θ, and without loss of generality we can take γ' to be this ray.

So, for sufficiently small δ, we have that (x_i) is within δ of some I-projection of a ray $\gamma' \in \Gamma_{\Theta'}$ adjacent to γ_Θ. $\square$

As noted at the end of the previous section, an immediate corollary of this result is that we can get the containment relation between faces and chambers by selecting a representative $\Theta \in \Delta_{(\mathbb{E},W)}$ in the interior of each face and chamber.

10.3. **Definability of a projective plane.** We now comment on extracting the set of rays corresponding to the projective plane using the definability results from the previous two sections.

Recall that we obtain the projective plane as follows. In the spherical building, we consider vertices labeled 1 to be projective points, and vertices labeled 2 to be projective lines. We use the incidence relation of being contained in a common chamber. Because we are working with a nice class of spaces, the theory of buildings tells us that we have a projective space over some field. So we can obtain a projective plane using the construction at the start of Proposition 8.3.1.

In the current setting, recall that each vertex of the spherical building corresponds to a singular ray γ. We will start with a vertex p_0 labeled 1, and select a vertex ℓ_0 labeled 2 which is not incident with p_0. This gives us a point p_0 and a line ℓ_0 not containing that point. Let $P(\ell_0)$ be the set of label 1 vertices which are incident with ℓ_0. These are the points on the line ℓ_0. For each such $p \in P(\ell_0)$, there is a label 2 vertex ℓ_{p,p_0} which is incident with both p and p_0. These are the lines which contain both points p and p_0. For each such ℓ_{p,p_0}, there is a set $P(\ell_{p,p_0})$ of label 1 vertices incident with ℓ_{p,p_0}. These are the points on the line ℓ_{p,p_0}. The projective plane is the collection of all points obtained this way. That is, it is the union

$$\bigcup_{p \in \ell_0} P(\ell_{p,p_0}).$$

In our structures, this will be realized as a definable subset of rays in $\Gamma_\star$ corresponding to some type 1 vertices in the spherical building at infinity.

The previous sections showed that we can quantify over the set $\Gamma_\star^{\Theta_k}$ of vertices of a given label Θ_k, and that given some $\gamma \in \Gamma_\star^{\Theta}$ corresponding to a known vertex with label Θ, we can quantify over rays corresponding to incident vertices of another given label Θ'.

We now use this to show definability of a projective plane across the asymptotic cones of a labeled symmetric space. We will also need to quantify over vertices of a given type which are simultaneously incident to two other vertices. This is not always possible, so we will select our constants carefully to enable us to do so in a special case.

Theorem 10.3.1. *Let M be a labeled symmetric space, where we have selected γ_1 and γ_2 to have the maximal possible distance in $D(\Theta_1, \Theta_2)$.*

Let $\mathcal{C}$ be the class of asymptotic cones $(N, d, \star, \gamma_1, \ldots, \gamma_K)$ of M.

Then the projective plane built from γ_1 and γ_2 is L-definable in the class $\mathcal{C}$.

Proof. Note that γ_1 and γ_2 in N continue to have the maximal possible distance.

We think of $\Gamma_\star^{\Theta_1}$ as a set of projective points and $\Gamma_\star^{\Theta_2}$ as a set of projective lines. So we are given a projective point γ_1 and a projective line γ_2 which is not incident with γ_1. The general construction of the projective plane and the definability results above provide the proof, once we have established that we can quantify over the set consisting of the projective lines $\gamma_2' \in \Gamma_\star^{\Theta_2}$ which are simultaneously incident with both the point γ_1 and some $\gamma_1' \in \Gamma_\star^{\Theta_1}$ which is incident with γ_2.

We will use the ultraproduct characterization of definability. We can work in the closure $\mathcal{C}'$ of $\mathcal{C}$ under ultraproducts, recalling that ultraproducts of euclidean buildings are still euclidean buildings with the induced atlas.

Denote the incidence relation of being contained in a common chamber by $\sim$. The set

$$P = \{\gamma_1' \in \Gamma_\star^{\Theta_1} : \gamma_1' \sim \gamma_2\}$$

of label 1 rays incident with γ_2 is definable by our earlier result. Our goal is to show the definability of the set A of rays $\gamma_2' \in \Gamma_{\Theta_2}$ which are simultaneously incident with γ_1 and some $\gamma_1' \in P$, that is

$$A = \{\gamma_2' \in \Gamma_\star^{\Theta_2} : \text{there is } \gamma_1' \in P \text{ such that } \gamma_2' \sim \gamma_1 \text{ and } \gamma_2' \sim \gamma_1'\}.$$

As usual, because A is a set of sequences, we check the definability of its finite projections.

We consider the zero set of a formula similar to the one we used for the proof of Theorem 10.2.1. Let ϕ be the formula given by

$$\inf_{\gamma_1' \in P} \inf_{(y_i) \in \pi_I(\Gamma_\star^{\Theta_2})} \left(\left(\sum_{i \in I} d(x_i, y_i) \right) + \sigma_{\gamma_1} + \sigma_{\gamma_1'} \right)$$

where σ_{γ_1} is the formula

$$\sum_{i \in I} \left(d(y_i, \gamma_1(i)) \approx i \cdot 2 \sin(\alpha/2) \right),$$

and similarly for $\sigma_{\gamma_1'}$ but with γ_1' in place of γ_1. Here, α is the minimal distance between elements of Γ_{Θ_1} and Γ_{Θ_2}.

Notice that $\pi_I(A)$ is the zero set of ϕ, for reasons similar to those in our proof of 10.2.1.

Now, suppose (x_i) is in the zero set of ϕ in an ultraproduct. Then (x_i) is the I-projection of some γ_2' which is incident with both γ_1 and some γ_1'. Because γ_1' is incident with γ_2, and because γ_2 is at the maximal possible distance from γ_1, we must have that γ_1 and γ_1' are distinct. So, there is a unique label 2 ray incident with both γ_1 and γ_1'.

Moreover, for similar reasons γ_1 and γ_1' are the ultralimits of a sequence of distinct pairs $\gamma_{1,j}$ and $\gamma_{1,j}'$ from each factor, each such pair having a unique shared label 2 incident ray $\gamma_{2,j}$. Thus γ_2' must be the unique ray arising as the ultralimit of the $\gamma_{2,j}$. So, (x_i) is in the ultraproduct of the zero sets. $\square$

The problem with generalizing the definability used in the argument above is that being not incident is not definable. In general, there are sequences of pairs from $\Gamma_\star^{\Theta_1} \times \Gamma_\star^{\Theta_2}$ which are always non-incident, but which tend toward an incident pair. Requiring that γ_1 and γ_2 are at the maximal distance avoids this and makes things easy to discuss in the ultraproduct.

10.4. **Field operations.** In this section, we show the existence of formulas whose zero sets are the graphs of the field operations constructed on a line in the projective plane, once we select some constants to set the scale of the field. Note that because we are working in a projective line, there is an infinity element ∞. We do not actually need the definability of the projective plane for this, since we will provide ourselves enough constants to carry out the construction. The constant ray γ_2 will serve as the projective line corresponding to the field.

The language L used below is taken to have constant symbols for $\gamma_1, \ldots \gamma_k$ and also the new constants we describe.

Theorem 10.4.1. *Let M be a labeled symmetric space with an additional 7 constant rays*

$$0, 1, \infty \in \Gamma_\star^{\Theta_1}(M)$$

$$\ell_0, \ell_1, \ell_\infty, \ell_\infty' \in \Gamma_\star^{\Theta_2}(M)$$

such that

- $0, 1, \infty$ *are distinct and are all incident with* γ_2*,*
- $\ell_0, \ell_1, \ell_\infty, \ell_\infty'$ *are distinct, not equal to* γ_2*, and satisfy*
  - ℓ_0 *is incident with* 0*,*
  - ℓ_1 *is incident with* 1*,*
  - $\ell_\infty, \ell_\infty'$ *are both incident with* ∞*,*
  - *there is no* $\gamma \in \Gamma_{\Theta_1}(M)$ *which is incident with all three of* $\ell_0, \ell_1, \ell_\infty$*.*

Let $\mathcal{C}$ be the class of asymptotic cones

$$(N, d, \star, \gamma_1, \ldots, \gamma_K, 0, 1, \infty, \ell_0, \ell_1, \ell_\infty, \ell_\infty')$$

of M.

Then the set $\mathbb{F}_\infty$ of rays $\gamma \in \Gamma_\star^{\Theta_1}$ incident with γ_2 is definable in $\mathcal{C}$, and there are L-formulas ϕ_+ and $\phi_\times$ such that for every $N \in \mathcal{C}$, the zero sets of ϕ_+ and $\phi_\times$ restricted to $\mathbb{F}_\infty$ are the graphs of the (extended) addition and multiplication operations of a field structure (with an infinite point) on $\mathbb{F}_\infty$.

That is, for $\gamma_a, \gamma_b, \gamma_c$ incident with γ_2, we have the following:

- $\phi_+(\gamma_a, \gamma_b, \gamma_c) = 0$ *iff* $\gamma_a + \gamma_b = \gamma_c$ *or* $(\gamma_a, \gamma_b) = (\infty, \infty)$,
- $\phi_\times(\gamma_a, \gamma_b, \gamma_c) = 0$ *iff* $\gamma_a \cdot \gamma_b = \gamma_c$ *or* $(\gamma_a, \gamma_b) \in \{(0, \infty), (\infty, 0)\}$.

Proof. The definability of the set of projective points incident with γ_2 was established in 10.2.1. This gets us the underlying set which we know can be identified with a projective line over a field $\mathbb{F}$.

For the formulas ϕ_+ and $\phi_\times$, consider the classic construction we discussed in Proposition 8.3.1 and note that triples on the graph can be described by a finite number of incidence relations and the constants we have given ourselves. For example, a triple (a, b, c) on the graph of addition (that is, satisfying $a + b = c$) can be characterized by the existence of several projective lines and points with a particular configuration of incidences.

So, we can construct the needed $\phi_+(x, y, z)$ or $\phi_\times(x, y, z)$ by writing a formula using infs over Γ_{Θ_1}, Γ_{Θ_2}, rays in Γ_{Θ_1} incident with a given ray in Γ_{Θ_2}, and rays in Γ_{Θ_2} incident with a given ray in Γ_{Θ_1}, and asserting that the rays have the required incidences by summing the formulas used in the incidence definability Theorem 10.2.1 which have the appropriate zero sets.

Such a formula will be zero exactly when a configuration of points and lines exists showing that (a, b, c) is on the graph. This completes the proof. $\square$

11. Instability of the Asymptotic Cones

In these sections, we use the results we have obtained to demonstrate that the order property holds in a uniform way in the class of asymptotic cones of labeled symmetric spaces defined over the reals. We will conclude that in this case, the spaces $(M, \frac{d}{n}, \star)$ approach the order property, and so by Proposition 4.0.9, there are many nonisomorphic asymptotic cones.

11.1. ${}^\rho\mathbb{R}_\mathcal{U}$ as a metric ultraproduct.

For the spaces we are interested in, the field in the asymptotic cone is the field ${}^\rho\mathbb{R}_\mathcal{U}$ below, originally defined by Abraham Robinson. A detailed discussion of this field can be found in chapters 3 and 4 of [16]. It is a nonarchimedean, real closed, valued field.

In this section we give the original definition of the field via discrete ultraproducts and valuations, and show how it can be constructed directly as a metric ultraproduct. We then easily observe that this structure is always unstable by exploiting the usual order for real closed fields in a bounded, noncompact subset. As a corollary, we reproduce the known result that when the continuum hypothesis fails, there are $2^{2^{\aleph_0}}$ many nonisomorphic such fields.

Definition 11.1.1. For $\mathcal{U}$ a nonprincipal ultrafilter on $\mathbb{N}$, we define ${}^\rho\mathbb{R}_\mathcal{U}$ to be the real closed valued field constructed as follows.

First denote the discrete ultraproduct of $\mathbb{R}$ with respect to $\mathcal{U}$ by

$$^*\mathbb{R}_{\mathcal{U}} = \prod_{\mathcal{U}} (\mathbb{R}, <, +, \cdot, 0, 1).$$

Let ρ be the equivalence class of $(e^{-n})_{n\in\mathbb{N}}$, which is a positive infinitesimal in $^*\mathbb{R}_{\mathcal{U}}$. Define the subring

$$M_0 = \{t \in {}^*\mathbb{R}_{\mathcal{U}} : |t| < \rho^{-k} \text{ for some } k \in \mathbb{N}\}$$

which can be informally thought of as the subset of "finite with respect to ρ" elements of $^*\mathbb{R}_{\mathcal{U}}$. Define

$$M_1 = \{t \in M_0 : |t| < \rho^{k} \text{ for all } k \in \mathbb{N}\}$$

which is the maximal ideal of non-units in M_0 and informally can be thought of as the set of "infinitesimal with respect to ρ" elements.

We define $^\rho\mathbb{R}_{\mathcal{U}}$ to be the field M_0/M_1.

The valuation v on $^\rho\mathbb{R}_{\mathcal{U}}$ is defined by $v(x) = \mathrm{st}(\log_\rho |x|)$. The associated norm $||x||_v$ is given by $e^{-v(x)}$, and the associated metric d_v is given by $d_v(x,y) = ||x-y||_v$.

Proposition 11.1.2. *The metric ultraproduct*

$$\prod_{\mathcal{U}} (\mathbb{R}, \sqrt[n]{d}, +, \cdot, 0, 1)_{n\in\mathbb{N}}$$

is isomorphic to

$$(^\rho\mathbb{R}_{\mathcal{U}}, d_v, +, \cdot, 0, 1)$$

where $\sqrt[n]{d}$ *is the metric given by*

$$\sqrt[n]{d}(x,y) = \sqrt[n]{d(x,y)}$$

for d *the usual metric on* $\mathbb{R}$*.*

Proof. We first need to establish that each $(\mathbb{R}, \sqrt[n]{d}, +, \cdot, 0, 1)$ is a metric structure in the same language. This requires checking that the suggested metric is in fact a metric, and that the functions $+, \cdot$ are controlled in a uniform way.

That $\sqrt[n]{d}$ is symmetric and positive definite is clear. The triangle inequality follows from subadditivity and monotonicity of $\sqrt[n]{x}$. So each $\sqrt[n]{d}$ is a metric.

Uniform continuity of $+$ with respect to $\sqrt[n]{d}$ does not depend on n nor on the magnitude of its inputs.

$$\begin{aligned}
\sqrt[n]{d}((x_1,y_1),(x_2,y_2)) \le \delta \quad &\iff \quad \max\left(\sqrt[n]{d}(x_1,x_2), \sqrt[n]{d}(y_1,y_2)\right) \le \delta \\
&\iff \quad \max\left(\sqrt[n]{d(x_1,x_2)}, \sqrt[n]{d(y_1,y_2)}\right) \le \delta \\
&\iff \quad \max\left(d(x_1,x_2), d(y_1,y_2)\right) \le \delta^n \\
&\implies \quad d(x_1+y_1, x_2+y_2) \le 2\delta^n \\
&\iff \quad \sqrt[n]{d}(x_1+y_1, x_2+y_2) \le \sqrt[n]{2}\delta \\
&\implies \quad \sqrt[n]{d}(x_1+y_1, x_2+y_2) \le 2\delta
\end{aligned}$$

Uniform continuity of $\cdot$ with respect to $\sqrt[n]{d}$ does not depend on n but does require restriction to bounded balls. The function Δ given by $\Delta(\varepsilon) = \min\{1, \frac{\varepsilon}{2R+1}\}$ can be used for the ball of radius R.

$$\begin{aligned}
\sqrt[n]{d}((x_1,y_1),(x_2,y_2)) \leq \delta \quad &\iff \quad \max\left(d(x_1,x_2), d(y_1,y_2)\right) \leq \delta^n \\
&\implies \quad d(x_1y_1, x_2y_2) \leq \delta^n ||x_1 + y_1|| + \delta^{2n} \\
&\implies \quad d(x_1y_1, x_2y_2) \leq \delta^n 2R + \delta^n \\
&\implies \quad d(x_1y_1, x_2y_2) \leq \delta^n (2R+1) \\
&\iff \quad \sqrt[n]{d}(x_1y_1, x_2y_2) \leq \delta(2R+1)^{1/n} \\
&\implies \quad \sqrt[n]{d}(x_1y_1, x_2y_2) \leq \delta(2R+1)
\end{aligned}$$

This verifies that each $(\mathbb{R}, \sqrt[n]{d}, +, \cdot, 0, 1)$ can be viewed as a metric structure in the same language $L = (+, \cdot, 0, 1)$.

We now describe the claimed isomorphism with $({}^\rho\mathbb{R}_\mathcal{U}, d_v, +, \cdot, 0, 1)$. Notice that we have the equality $\sqrt[n]{x} = e^{\log(x)/n}$ for nonzero x. Moreover, the ultralimit of the function $e^{\log(x)/n}$ is the function $e^{-\log_\rho(x)}$ by definition of ρ. We can take this to hold for 0 as well by the convention that $e^{-\infty}$ is 0 when defining the valuation norm. Thus, the sequences $(a_n)_{n\in\mathbb{N}}$ with $\sqrt[n]{d}(0, a_n)$ having bounded ultralimit are exactly those sequences t of ${}^*\mathbb{R}_\mathcal{U}$ with $\rho^k < t < \rho^{-k}$ for some $k \in \mathbb{N}$, i.e. the elements $t \in^\rho \mathbb{R}_\mathcal{U}$. This establishes that the underlying set of sequences for both structures is the same. It is also clear that the distance between two sequences (a_n) and (b_n) is the same in either metric, and so the identity function is an isometry between the structures which fixes 0 and 1. Since the operations in both structures are defined in the same way, this is an isomorphism. □

Proposition 11.1.3. $({}^\rho\mathbb{R}_\mathcal{U}, d_v, +, \cdot, 0, 1)$ *is unstable.*

Proof. Let $\phi(x, y)$ be the formula

$$\inf_z]_1^2 d(x + (z \cdot z), y).$$

The field contains a copy of the naturals, so for each $n \in \mathbb{N}$, let $a_n = n \in^\rho \mathbb{R}_\mathcal{U}$. Note that the distance between any distinct $m, n \in \mathbb{N}$ is

$$\lim_\mathcal{U} \sqrt[n]{|m-n|} = 1.$$

So the sequence (a_n) is 1-separated and contained in the closed ball of radius 1 at the origin 0.

Whenever $m \leq n$, we have

$$\inf_z]_1^2 d(m + z^2, n) = 0$$

since there is some $z \in \bar{B}_1$ with $m + z^2 = n$.

Whenever $m > n$, we have

$$\inf_z]_1^2 d(m + z^2, n) = 1$$

which is witnessed by $z = 0$.

So ϕ and (a_n) demonstrate the order property. □

Corollary 11.1.4. *The sequence of structures* $(\mathbb{R}, \sqrt[n]{d}, +, \cdot, 0, 1)$ *approaches the order property.*

Proof. This is clear by using the formula ϕ above and the copy of the naturals in each factor. □

Notice that each $(\mathbb{R}, \sqrt[n]{d}, +, \cdot, 0, 1)$ is stable because it is proper. The instability arises because of the distortion of the metric. Compare this to our earlier observation that asymptotic cones of $(\mathbb{R}, d)$ are just isomorphic to $(\mathbb{R}, d)$ again.

We can easily conclude the following known result (for example, see the discussion after Theorem 1.8 of [15]).

Corollary 11.1.5. *If the continuum hypothesis holds, then* ${}^\rho\mathbb{R}_\mathcal{U}$ *is always saturated. If the continuum hypothesis fails, there are* $2^{2^{\aleph_0}}$ *many nonisomorphic* ${}^\rho\mathbb{R}_\mathcal{U}$ *as we vary* $\mathcal{U}$*.*

Proof. The saturation claim follows from the construction of ${}^\rho\mathbb{R}_\mathcal{U}$ as a countable ultraproduct. The many models claim follows from applying 4.0.9 and the previous corollary. □

11.2. **Approaching the order property in the projective line.** When our symmetric space M is given by a connected, absolutely simple (i.e. simple over $\mathbb{C}$) real Lie group G, the field associated to the spherical building at infinity for M is just the real field $\mathbb{R}$. But, in this case, the field associated to the spherical building of the $\mathcal{U}$ asymptotic cone is known to be ${}^\rho\mathbb{R}_\mathcal{U}$ (see Theorem 2.6 and the surrounding discussion of [15]). Each of these is a real closed field, and so we can use an analog of the usual construction of the ordering for real closed fields to find ordered sequences.

Note that these fields are appearing as subsets of rays in our structures. Thus, their metrics are not the standard ones, since when viewed as subsets of rays we have only given them the (bounded) metric induced from $\angle_\infty$. A different way to carry out our argument would be to focus on the spherical building structure, noting that the our earlier definability results are independent of the choice of ultrafilter, and so imply definability of the spherical building and field in the symmetric spaces involved in the asymptotic cone. In that view, one could see the spherical building of the asymptotic cone as the ultraproduct of spherical buildings of the symmetric spaces, and connect this observation to the instability shown in the previous section. Here, we will just directly check that our structures approach the order property.

We are restricting ourselves to the case where G is defined over $\mathbb{R}$, but suspect that the methods developed in this thesis might help in understanding the complex case where one would obtain a copy of the analogous field ${}^\rho\mathbb{C}_\mathcal{U}$ with metric given by $\angle_\infty$. The appearance of this field and metric in the structure may help in understanding such asymptotic cones.

We work in the language L with $K+7$ many constants for rays $\gamma_1, \ldots, \gamma_K$ and $0, 1, \infty$, $\ell_0, \ell_1, \ell_\infty, \ell'_\infty$. Write Consts for the set of these constants.

Theorem 11.2.1. *Let M be as in Theorem 10.4.1, and moreover assume its underlying metric space is given by G/K for G some connected, absolutely simple Lie group defined over $\mathbb{R}$.*

For each $n \in \mathbb{N}_+$, let M_n be the n-th factor

$$M_n = (M, \frac{d}{n}, \star, n\gamma : \gamma \in \mathrm{Consts})$$

in the construction of asymptotic cones of M.

Then the sequence of L-structures $(M_n : n \in \mathbb{N}_+)$ approaches the order property.

Proof. Each M_n is a symmetric space, so there is a spherical building at infinity where we can (externally to the logic) construct the real field on the projective line corresponding to γ_2. Consider the sequence $(\nu_k : k \in \mathbb{N})$ of rays ν_k corresponding to the natural numbers k constructed in (M, d). We will use this sequence to demonstrate that we approach the order property.

Because the incidence relation is defined via chambers in the building at infinity, scaling the metric on M does not affect the incidence. So each of these rays

$$\nu_k \in \Gamma_\star^{\Theta_1}(M)$$

has an n-th scaling

$$n\nu_k \in \Gamma_\star^{\Theta_1}(M_n)$$

which represents the natural k in the field constructed in M_n. Moreover, in any asymptotic cone, we obtain the elements $(n\nu_k : n \in \mathbb{N}_+)/\sim$, which for $k \in \mathbb{N}$ represent the natural k in the field constructed in the asymptotic cone. We will just write these as the corresponding natural for convenience. That is, we write $k = (n\nu_k : n \in \mathbb{N}_+)/\sim$ for each $k \in \mathbb{N}$.

Now, because the original naturals and ∞ in M were all distinct rays and we are just using rescalings, their ultralimits in the asymptotic cone must pairwise span flat sectors (i.e. diverge linearly, not piecewise linearly). That is, in any asymptotic cone, the naturals form a discrete set with a lower bound δ on the pairwise distances, independently of the choice of ultrafilter. The bound δ is distance corresponding to the minimal nonzero element of $D(\Theta_1, \Theta_1)$.

Let $\mathcal{C}$ be the class of asymptotic cones

$$(N, d, \star, \gamma_1, \ldots, \gamma_K, 0, 1, \infty, \ell_0, \ell_1, \ell_\infty, \ell'_\infty)$$

of M.

Recall that we can quantify over the set $\mathbb{F}$ of rays incident with γ_2, and that this set of points corresponds to the underlying set of the field we know to be ${}^\rho\mathbb{R}_\mathcal{U}$. Note that we are only saying this is a copy of ${}^\rho\mathbb{R}_\mathcal{U}$ as a field. The metric in this case is inherited from rays and is bounded, not the standard metric constructed in the previous section.

Consider the formula $\psi_{\leq}(x,y)$ given by

$$\inf_{t\in\mathbb{F}}\inf_{s\in\mathbb{F}}\left(\phi_+(x,s,y)+\phi_\times(t,t,s)\right).$$

By $\aleph_1$-saturation of the ultraproduct, we have that $\psi_{\leq}(x,y)=0$ exactly when there are t,s such that $t^2=s$ and $x+s=y$; i.e. exactly when x plus a square is y. So with the sequence $(a_k)_{k\in\mathbb{N}}=(0,1,2,\dots)$ we have that $\psi_{\leq}(a_i,a_j)=0$ iff $i\leq j$.

We now check that $(M_n : n\in\mathbb{N}_+)$ approaches the order property using this last paragraph. This mostly comes down to using that the above paragraph applies independently of the choice of ultrafilter. Our candidate sequence to order in each M_n will be the sequence $(n\nu_k : k\in\mathbb{N})$ of naturals in that space.

First, if there were an $\varepsilon>0$ and $k\in\mathbb{N}$ with infinitely many n where in M_n, there were a_i and a_j in

$$(a_0,\dots,a_k)=(n\nu_0,\dots,n\nu_k)$$

with $i\leq j$ but with $\psi_{\leq}(a_i,a_j)>\varepsilon$, then in some ultraproduct we would get a pair a,b on the field line representing naturals with $a\leq b$ but with $\psi_{\leq}(a,b)\geq\varepsilon>0$. This is impossible. So it must be that for all $\varepsilon>0$ and $k\in\mathbb{N}$, there are cofinitely many n where in M_n we have $\psi_{\leq}(a_i,a_j)\leq\varepsilon$ whenever $i\leq j<k$.

Next, suppose it were possible to find for every $\varepsilon>0$ infinitely many n where in M_n there were a_i and a_j in

$$(a_0,\dots,a_k)=(n\nu_0,\dots,n\nu_k)$$

with $j<i$ but with $\psi_{\leq}(a_i,a_j)\leq\varepsilon$. Then in some ultraproduct we would get a pair a,b on the field line satisfying representing naturals with $a>b$ but with $\psi_{\leq}(a,b)=0$. Again, this is impossible. So there must be some fixed $\varepsilon^*=\varepsilon>0$ for which all but possibly finitely many M_n have $\psi_{\leq}(a_i,a_j)>\varepsilon$ whenever $k\geq i>j$.

Combining the two previous paragraphs shows that M_n approaches the order property, witnessed by the formula $\psi_{\leq}$ and the tuples $(n\nu_0,\dots n\nu_k)$ for each $k\in\mathbb{N}$. $\square$

We immediately have the following. The second part reproduces the result in [15].

Corollary 11.2.2. *All asymptotic cones of M are unstable in the language with no symbols other than d and $\star$. Moreover, if the continuum hypothesis fails, there are $2^{2^{\aleph_0}}$ many asymptotic cones of M.*

Proof. The constants we used above can be used as parameters instead to get an empty language. The naturals are 1-separated and ordered by $\phi_{\leq}$ in any of the asymptotic cones. The other part of the claim follows by applying 4.0.9 to the sequence $(M_n : n\in\mathbb{N}_+)$. $\square$

References

[1] Jean-Martin Albert. *Strong conceptual completeness and various stability-theoretic results in continuous model theory.* PhD thesis, McMaster University, 2011.

[2] Werner Ballman, Mikhael Gromov, and Viktor Schroeder. *Manifolds of Nonpositive Curvature.* Birkhäuser, 1985.

[3] I. Ben Yaacov, A. Berenstein, C. W. Henson, and A. Usvyatsov. Model theory for metric structures. *Model Theory with Applications to Algebra and Analysis, Vol. II*, London Math. Soc. Lecture Notes Series, no. 350:315–427, 2008.

[4] Itaï Ben Yaacov. Continuous first order logic for unbounded metric structures. (October):1–21, 2009.

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

[6] C. D. Bennet, P. N. Schwer, and K. Struyve. On axiomatic definitions of non-discrete affine buildings. *Advances in Geometry*, 14, 2014.

[7] Armand Borel and Lizhen Ji. *Compactifications of Symmetric and Locally Symmetric Spaces.* Birkhäuser, 2006.

[8] Martin Bridson and Andre Haefliger. *Metric spaces of non-positive curvature*, volume 319. Springer, 1999.

[9] Patrick Eberlein. Structure of manifolds of nonpositive curvature. In *Global Differential Geometry and Global Analysis 1984*, volume 1156 of *Lecture Notes in Mathematics*, pages 86–153. Springer, 1985.

[10] Patrick Eberlein. *Geometry of Nonpositively Curved Manifolds.* Chicago Lectures in Mathematics, 1997.

[11] I Farah and S Shelah. A dichotomy for the number of ultrapowers. *Journal of Mathematical Logic*, 10:45–81, 2010.

[12] Ilijas Farah, Bradd Hart, and David Sherman. Model theory of operator algebras ii: model theory. *Israel Journal of Mathematics*, 201(1):477–505, 2014.

[13] M Gromov. Groups of polynomial growth and expanding maps. *Publications Mathématiques de l'IHÉS*, 1981.

[14] Bruce Kleiner and Bernhard Leeb. Rigidity of quasi-isometries for symmetric spaces and euclidean buildings. *Comptes Rendus de l'Académie des Sciences-Series I-Mathematics*, 324(6):639–643, 1997.

[15] Linus Kramer, Saharon Shelah, Katrin Tent, and Simon Thomas. Asymptotic cones of finitely presented groups. *Advances in Mathematics*, 193(1):142–173, May 2005.

[16] A. H. Lightstone and A. Robinson. Nonarchimedean fields and asymptotic expansions. 1975.

[17] Anne Parreau. *Dégénérescences de sous-groupes discrets de groupes de Lie semisimples et actions de groupes sur des immeubles affines.* PhD thesis, Université Paris-Sud, Orsay, 2000.

[18] Betti Salzberg. Buildings and shadows. *Aequationes mathematicae*, 25:1–20, 1982.

[19] L. van den Dries and A. J. Wilkie. Gromov's theorem on groups of polynomial growth and elementary logic. *J. ALGEBRA.*, 374:349–374, 1984.

[20] Veblen and Young. *Projective Geometry*, volume 1. Boston Ginn, 1910.