Lectures on quasi-isometric rigidity

Michael Kapovich

June 28, 2012

1   Introduction

What is Geometric Group Theory (GGT)?

Historically (19th century), groups appeared as automorphism groups of some structures:

- Polynomials (field extensions) — Galois groups.

- Vector spaces, possibly equipped with a bilinear form— Matrix groups.

- Complex analysis, complex ODEs — Monodromy groups.

- Partial differential equations — Lie groups (possibly infinite-dimensional ones)

- Various geometries — Isometry groups of metric spaces, both discrete and nondiscrete.

Goal of GGT is to study finitely-generated (f.g.) groups G as automorphism groups (symmetry groups) of metric spaces X.

Central question: How algebraic properties of a group G reflect in geometric properties of a metric space X and, conversely, how geometry of X reflects in algebraic structure of G.

This interaction between groups and geometry is a fruitful 2-way road.

Inspiration: Simple noncompact connected Lie groups — Irreducible symmetric spaces of noncompact type (E.Cartan et al).

Here there is an essentially 1-1 correspondence between algebraic objects (a Lie group of a certain type) and geometric objects (certain symmetric spaces). Namely, given a Lie group G on constructs a symmetric space $X = G/K$ (K is a maximal compact subgroup of G) and, conversely, every symmetric space corresponds to a Lie group G (its isometry group) and this group is unique.

Imitating this correspondence is an (unreachable) goal for GGT.

1

2 Cayley graphs and other metric spaces

Recall that we are looking for a correspondence:

$$\text{groups} \longleftrightarrow \text{metric spaces}$$

The first step is to associate with a f.g. group G a metric space X. Let G be a group with a finite generating set $S = \{s_1, ..., s_k\}$. It is sometimes convenient to assume that S is symmetric, i.e., $\forall s \in S$, $s^{-1} \in S$. Then we construct a graph X, where the vertex set $V(X)$ is the group G itself and edges of X are

$$[g, gs_i], s_i \in S, g \in G.$$

(If $gs_i = gs_j$, i.e., $s_i = s_j$, then we treat $[g, gs_i], [g, s_j]$ as distinct edges, but this is not very important.) We do not orient edges.

The resulting graph X is called *Cayley graph* of the group G with respect to the generating set S.

For examples of Cayley graphs — see the 1st discussion section.

Then the group G acts (by multiplication on the left) on X: Every $\gamma \in G$ defines a map

$$\gamma(g) = \gamma g, \quad g \in G.$$

Clearly, edges are preserved. Since S is a generating set of G, the graph X is connected.

Metric on a graph: If X is any connected graph, then we declare every edge of X to have unit length. Then we have a well-defined notion of length of a path in X. Distance between points in X is the length of the shortest path in X connecting these points.

Exercise 2.1. *Shortest path always exists.*

If you like, you can think of connecting only vertices of X by paths, as we will see, this is not very important. The restriction of this metric to G is called a *word-metric* on G. Here is why:

Example 2.2. *Let X be Cayley graph of a group G. Distance $d(1, g)$ from $1 \in G$ to $g \in G$ is the same thing as the "norm" $|g|$ of g, the minimal number m of symbols in the decomposition (a "word in the alphabet $S \cup S^{-1}$")*

$$g = s_{i_1}^{\pm 1} s_{i_2}^{\pm 1} ... s_{i_m}^{\pm 1}$$

of g as a product of generators and their inverses. Note: If $g = 1$ then $m = 0$.

Other metric spaces which appear naturally in GGT: Connected Riemannian manifolds (M, ds^2). In this case, the distance between points is

$$d(x, y) = \inf\{\int_p ds = \int_0^T |p'(t)| dt\}$$

where infimum is taken over all paths p connecting x to y.

This, we have a correspondence: **Groups** $\rightarrow$ **Metric spaces**,

$$\text{Cayley} : G \rightarrow X = \text{Cayley graph of } G.$$

Is this the only correspondence? (No!) Is this map "Cayley" well defined ? (Not really, since G has many generating sets.)

Definition 2.3. *Let X be a metric space and G is a group acting on X. The action $G \curvearrowright X$ is called* geometric *if:*

1. G acts isometrically on X.

2. G acts properly discontinuously on X (i.e., $\forall$ compact $C \subset X$, the set $\{g \in G : gC \cap C \neq \emptyset\}$ is finite).

3. G acts cocompactly on X: X/G is compact.

Informally, a group X is a group of (discrete) symmetries of X if G acts geometrically on X.

Example 2.4. *G is a f.g. group, X is its Cayley graph. Question: What is the quotient graph X/G?*

Example 2.5. *M is a compact connected Riemanninan manifold, $\pi = \pi_1(M)$, the fundamental group, $X = \tilde{M}$ is the universal cover of M (with lifted Riemannian metric), π acts on X as the group of covering transformations for the covering $X \rightarrow M$.*

More generally, let $\phi : \pi \rightarrow G$ be an epimorphism, $X \rightarrow M$ be the covering corresponding to $Ker(\phi)$. Then the group of covering transformations of $X \rightarrow M$ is G and G acts geometrically on X.

Note: For every f.g. group G there exists a compact Riemannian manifold M (in every dimension ≥ 2) with an epimorphism $\pi_1(M) \rightarrow G$. Thus, we get another correspondence Groups $\rightarrow$ metric spaces,

$$\text{Riemann} : G \rightarrow X = \text{a covering space of some } M \text{ as above}.$$

Thus, we have a problem on our hands, we have too many candidates for the correspondence **Groups** $\rightarrow$ **Spaces** and these correspondences are not well-defined. What do different spaces on which G acts geometrically have in common?

3 Quasi-isometries

Definition 3.1. *a. Let X, X' be metric spaces. A map $f : X \rightarrow X'$ is called (L, A)-quasi-isometry if:*

1. f is (L, A)-Lipschitz:

$$d(f(x), f(y)) \leq Ld(x, y) + A.$$

2. There exists an (L, A)-Lipschitz map $\bar{f} : X' \to X$, which is "quasi-inverse" to f:

$$d(\bar{f}f(x), x) \leq A, \quad d(f\bar{f}(x'), x') \leq A.$$

b. Spaces X, X' are quasi-isometric to each other if there exists a quasi-isometry $X \to X'$.

Note, if $A = 0$ then such f is a bilipschitz homeomorphism; if $L = 1, A = 0$ then f is an isometry.

Example 3.2. *1. Every bounded metric space is QI to a point.*

2. $\mathbb{R}$ is QI to $\mathbb{Z}$.

3. Every metric space is QI to its metric completion.

Here and in what follows I will abbreviate "quasi-isometry" and "quasi-isometric" to QI.

Exercise 3.3. • *Every quasi-isometry is "quasi-surjective":*
  $\exists C < \infty | \forall x' \in X', \exists x \in X | d(x', f(x)) \leq C.$

  • *Show that a map $f : X \to X'$ is a quasi-isometry iff it is quasi-surjective and is a "quasi-isometric embedding": $\exists L, \exists A$ so that $\forall x, y \in X$:*

$$\frac{1}{L}d(x, y) - A \leq d(f(x), f(y)) \leq Ld(x, y) + A.$$

  • *Composition of quasi-isometries is again a quasi-isometry.*

  • *Quasi-isometry of metric spaces is an equivalence relation.*

Exercise 3.4. *1. Let S, S' be two finite generating sets for a group G and d, d' be the corresponding word metrics. Then the identity map $(G, d) \to (G, d')$ is an $(L, 0)$-quasi-isometry for some L.*

2. G is QI to its Cayley graph X. The map $G \to X$ is the identity. What is the quasi-inverse?

Lemma 3.5. *(Milnor–Schwartz lemma). Suppose that G acts geometrically on a "nice" metric space X (e.g. a graph or a Riemannian manifold). Then G is f.g. and $(\forall x \in X)$ the orbit map $g \mapsto g(x), G \to X$, is a quasi-isometry, where G is equipped with word-metric.*

Thus, if instead of isometry classes of metric spaces, we use their QI classes, then both *Cayley* and *Riemann* correspondences are well-defined and are equal to each other! Now, we have a well-defined map

$$geo : \text{f.g. groups} \longrightarrow \text{QI equivalence classes of metric spaces.}$$

Problem: This map is very far from being 1-1, so our challenge is to "estimate" the fibers of this map.

Exercise 3.6. *Show that half-line is not QI to any Cayley graph. Hint: Prove first that every unbounded Cayley graph contains an isometrically embedded copy of* $\mathbb{R}$. *Use Arzela-Ascoli theorem.*

Example 3.7. *Every finite group is QI to the trivial group.*

In particular, from the QI viewpoint, the entire theory of finite groups (with its 150 year-old history culminating in the classification of finite simple groups) becomes trivial. Is this good news or is this bad news?

This does not sound too good if we were to recover a group from its geometry (up to an isomorphism). Is there a natural equivalence relation on groups which can help us here?

4 Virtual isomorphisms and QI rigidity problem

In view of M-S lemma, the following provide examples of quasi-isometric groups:

1. If $G' < G$ is a finite-index subgroups then G is QI to G'. (G' acts on G isometrically and faithfully so that the quotient is a finite set.)

2. If $G' = G/F$, where F is a finite group, then G is QI to G'. (G acts isometrically and transitively on G' so that the action has finite kernel.)

Combining these two examples we obtain

Definition 4.1. *1. G_1 is VI to G_2 if there exist finite index subgroups $H_i \subset G_i$ and finite normal subgroups $F_i \triangleleft H_i, i = 1, 2$, so that the quotients H_1/F_1 and H_2/F_2 are isomorphic.*

2. A group G is said to be virtually *cyclic if it is VI to a cyclic group. Similarly, one defines virtually abelian groups, virtually free groups, etc.*

Lemma 4.2. *VI is an equivalence relation.*

By M-S lemma,
$$VI \Rightarrow QI$$

Thus, if we were to recover groups from their geometry (treated up to QI), then the best we can hope for is to recover a group up to VI. This is bad news for people in the finite group theory, but good news for the rest of us.

Remark 4.3. *There are some deep and interesting connections between theory of finite group and GGT, but quasi-isometries do not see these.*

Informally, *quasi-isometric rigidity* is the situation when the arrow $VI \Rightarrow QI$ can be reversed.

Definition 4.4. *1. We say that a group G is QI rigid if every group G' which is QI to G, is in fact VI to G.*

2. We say that a class C of group is QI rigid if every group G' which is QI to some $G \in C$, there exists $G'' \in C$ so that G' is VI to G''.

3. A property P of groups is said to be "geometric" or "QI invariant" whenever the class of groups satisfying P is QI rigid.

Note that studying QI rigidity and QI invariants is by no means the only topic of GGT, but this will be the topic of my lectures.

5 Examples and non-examples of QI rigidity

At the first glance, any time QI rigidity holds (in any form), it is a minor miracle: How on earth are we supposed to recover precise algebraic information from something as sloppy as a quasi-isometry? Nevertheless, instances of QI rigidity abound.

Examples of QI rigid groups/classes/properties (all my groups are finitely-generated, of course):

- Free groups.

- Free abelian groups.

- Class of nilpotent groups.

- Class of fundamental groups of closed (compact, without boundary) surfaces.

- Class of fundamental groups of closed (compact, without boundary) 3-dimensional manifolds.

- Class of finitely-presentable groups.

- Class of hyperbolic groups.

- Class of amenable groups.

- Class of fundamental groups of closed n-dimensional hyperbolic manifolds. For $n \geq 3$ this result, due to P.Tukia, will be the central theorem of my lectures.

- Class of discrete cocompact subgroups Γ in a simple noncompact Lie group G.

- Every discrete subgroup Γ in a simple noncompact Lie group G so that G/Γ has finite volume. For instance, every group which is QI to $SL(n, \mathbb{Z})$ is in fact VI to $SL(n, \mathbb{Z})$.

- Solvability of the word problem (say, for finitely-presented groups).

- Cohomological dimension over $\mathbb{Q}$.

- Admitting a "geometric" action on a contractible CW-complex (i.e., an action which is cocompact on each skeleton is cocompact and properly discontinuous).

- Admitting an amalgam decomposition (amalgamated free product or HNN decomposition) over a finite subgroup.

- Admitting an amalgam decomposition over a virtually cyclic subgroup.

Rule of thumb: The closer a group (a class of groups) is to a Lie group, the higher are the odds of QI rigidity.

Examples of failure of QI rigidity:

- Suppose that S is a closed surface of genus ≥ 2 and $\pi = \pi_1(S)$. Then $\mathbb{Z} \times \pi$ is QI to any Γ which appears in any central extension

$$1 \to \mathbb{Z} \to \Gamma \to \pi \to 1.$$

  For instance, the fundamental group Γ of the unit tangent bundle to Σ appears like this.

- In particular, the property of being the fundamental group of a compact nonpositively-curved Riemannian manifold with convex boundary is not QI invariant.

- There are countably many VI classes of groups which act geometrically on the hyperbolic 3-space. All these groups are QI to each other by M-S lemma. Same for all irreducible nonpositively curved symmetric spaces of dimension ≥ 3.

- Class of solvable groups is not QI rigid.

- Class of simple groups is not QI rigid.

- Class of residually-finite groups is not QI rigid.

- Property T is not QI invariant.

Few open problems:

- Is the class of π_1 of a closed aspherical n-dimensional orbifolds QI rigid?

- Is the class of polyciclic groups QI rigid (conjecturally, yes).

- Is the class of Right-Angled Artin Groups QI rigid?

- Are random finitely-presented groups QI rigid?

Where do the tools of GGT come from? Almost everywhere! Here are some examples:

- Group theory (of course)

- Geometry (of course)

- Topology (point-set topology, geometric topology, algebraic topology)

- Lie theory

- Analysis (including PDEs, functional analysis, real analysis, complex analysis, etc.)

- Probability

- Logic

- Dynamical systems

- Homological algebra

- Combinatorics

In these lectures, I will introduce two (of many) tools of QI rigidity: Ultralimits (coming from logic) and Quasiconformal maps (geometric analysis and real analysis).

6 Ultralimits and asymptotic cones

Motivation: Quasi-isometries are not nice maps, they need not be continuous, etc. We will use ultralimits to convert quasi-isometries to homeomorphisms.

Definition 6.1. *An* ultrafilter *on the set $\mathbb{N}$ of natural numbers is a finitely-additive (σ-additive) measure defined for all subsets of $\mathbb{N}$ and taking only the values 0 and 1.*

In other words, $\omega : 2^{\mathbb{N}} \to \{0, 1\}$ is:

- *Finitely-additive: $\omega(A \cup B) = \omega(A) + \omega(B) - \omega(A \cap B)$.*

- *$\omega(\emptyset) = 0$.*

We will say that a subset E of $\mathbb{N}$ is ω-large if $\omega(E) = 1$. Similarly, we will say that a property $P(n)$ holds for ω-all natural numbers if $\omega(\{n : P(n)$ is true $\}) = 1$.

Trivial examples of ultrafilters are such that $\omega(\{n\}) = 1$ for some $n \in \mathbb{N}$. I will always assume that ω vanishes on all finite sets.

Existence of ultrafilters does not follow from the Zermelo-Frankel (ZF) axioms of set theory, but follows from ZFC (ZF+ Axiom of Choice).

We will use ultrafilters to define limits of sequences:

Definition 6.2. *Let X be a Hausdorff topological space and ω is an ultrafilter (on $\mathbb{N}$). Then, for a sequence (x_n) of points $x_n \in X$, we define ω-limit (ultralimit), $\lim_\omega x_n = a$, by:*

For every neighborhood U of a the set $\{n \in \mathbb{N} : x_n \in U\}$ is ω-large.

In other words, $x_n \in U$ for ω-all n.

As X is assumed to be Hausdorff, $\lim_\omega x_n$ is unique (if it exists).

Exercise 6.3. *If $\lim x_n = a$ (in the usual sense) then $\lim_\omega x_n = a$ for every ω.*

I will fix an ultrafilter ω once and for all.

Lemma 6.4. *If X is compact then every sequence in X has ultralimit.*

In particular, every sequence $t_n \in \mathbb{R}_+$ has ultralimit in $[0, \infty]$.

Exercise 6.5. *What is the ultralimit of the sequence $(-1)^n$ in $[0, 1]$?*

Our next goal is to define ultralimit for a sequence of metric spaces (X_n, d_n). The definition is similar to Cauchy completion of a metric space: Elements of the ultralimit will be equivalence classes of sequences $x_n \in X_n$. For every two sequences $x_n \in X_n, y_n \in Y_n$ we define

$$d_\omega((x_n), (y_n)) := \lim_\omega d_n(x_n, y_n) \in [0, \infty].$$

Exercise 6.6. *Verify that d_ω is a pseudo-metric. (Use the usual convention $\infty + a = \infty$, for every $a \in \mathbb{R} \cup \{\infty\}$.)*

Of course, some sequences will be within zero distance from each other. As in the definition of Cauchy completion, we will identify such sequences (this is our equivalence relation). After that, d_ω is "almost" a metric: The minor problem is that sometimes d_ω is infinite. To handle this problem, we introduce a sequence of "observers", points $p_n \in X_n$. Then, we define $\lim_\omega X_n = X_\omega$, the ultralimit of the sequence of "pointed" metric spaces (X_n, p_n) to be the set of equivalence classes of sequences $x_n \in X_n$ so that

$$\lim_\omega d_n(x_n, p_n) < \infty.$$

Informally, X_ω consists of equivalence classes of sequences which the "observers" can see.

In case $(X_n, d_n) = (X, d)$, we will refer to $\lim_\omega X_n$ as a *constant* ultralimit.

9

Exercise 6.7.   • *If X is compact then the constant ultralimit $\lim_\omega X$ is homeomorphic to X (for any sequence of observers).*

   • *Suppose that X admits a geometric group action. Then the constant ultralimit $\lim_\omega X$ does not depend on the choice of the observers.*

   • *Suppose that X is a metric space where every closed metric ball is compact. Then for every bounded sequence $p_n \in X$ the constant ultralimit $\lim_\omega X$ is homeomorphic to X.*

   • *For instance, $\lim_\omega \mathbb{R}^k = \mathbb{R}^k$.*

The ultralimits that we will use are not quite constant: Take a metric space (X, d) and a sequence of positive scale factors λ_n so that $\lim_\omega \lambda_n = 0$. Then take $d_n := \lambda_n d$. Hence, the sequence (X, d_n) consists of rescaled copies of (X, d).

Definition 6.8. *An asymptotic cone of X, denoted $Cone(X)$ is the ultralimit of the sequence of pointed metric spaces: $Cone(X) = \lim_\omega (X_n, \lambda_n d, p_n)$.*

Note that, in general, the asymptotic cone depends on the choices of ω, (λ_n) and (p_n).

Exercise 6.9. *Let $G = \mathbb{Z}^n$ be the free abelian group with its standard set of generators. Let $X = G$ with the word metric. Then $Cone(X)$ is isometric to $\mathbb{R}^n$ with the ℓ_1-metric corresponding to the norm*

$$\|(x_1, ..., x_n)\| = |x_1| + ... + |x_n|.$$

Lemma 6.10. *Suppose that X is the hyperbolic space $\mathbb{H}^k$, $k \geq 2$. Then $X_\omega = Cone(X)$ (for all choices) is a tree which branches at every point and has infinite (continual) degree of branching at every point x_ω: The cardinality of the number of components of $X_\omega - \{x_\omega\}$ is continuum.*

Quasi-isometries and asymptotic cones. Suppose that $f : X \to X'$ is an (L, A)-quasi-isometric embedding:

$$\frac{1}{L} d(x, y) - A \leq d(f(x), f(y)) \leq L d(x, y) + A.$$

Pick a sequence of scale factors λ_n a sequence of observers $p_n \in X$ and their images $q_n = f(x_n)$. Then, after rescaling the metrics on X, X' by λ_n, we obtain:

$$\frac{1}{L} d(x, y) - \lambda_n A \leq d(f(x), f(y)) \leq L d(x, y) + \lambda_n A.$$

Thus, after taking the ultralimit:

$$f_\omega : X_\omega \to X'_\omega, \quad f_\omega((x_n)) = (f(x_n)).$$

we get:

$$\frac{1}{L}d_\omega(x,y) \le d_\omega(f_\omega(x), f_\omega(y)) \le L d_\omega(x,y)$$

for all $x, y \in X_\omega$. Thus, f_ω is a bilipschitz embedding, since the additive constant A is gone! Even better, if f was quasi-surjective, then f_ω is surjective. Thus, $f_\omega : X_\omega \to X'_\omega$ is a homeomorphism!

The same principle applies to sequences of quasi-isometric embeddings/quasi-isometries as longs as the constants L, A are fixed.

Exercise 6.11. $\mathbb{R}^n$ *is QI to* $\mathbb{R}^m$ *iff* $n = m$.

Exercise 6.12. *Suppose that* $\mathbb{R}^n \to \mathbb{R}^n$ *is a QI embedding. Then f is quasi-surjective. Hint: If not, then, taking an appropriate sequence of scaling factors and observers, and passing to asymptotic cones, we get* $f_\omega : \mathbb{R}^n \to \mathbb{R}^n$, *a bilipschitz embedding which is not onto. This map has to be open (since dimensions of domain and range are the same), it is also proper since it is bilipschitz. Thus, f_ω is also closed. Hence f_ω is onto.*

Unfortunately, we cannot tell $\mathbb{H}^n$ from $\mathbb{H}^m$ using asymptotic cones since all cones are isometric to the same tree!

Morse Lemma. Let $X = \mathbb{H}^n$ be a hyperbolic space. A *quasi-geodesic* in X is a QI embedding $f : I \to X$, where I is an interval in $\mathbb{R}$ (either finite or infinite).

Lemma 6.13. *There exists a function $D(L, A)$ so that every (L, A)-quasi-geodesic in X is D-close to a geodesic.*

Proof. Quasi-geodesics in X yield bi-Lipschitz embedded curves in the tree $Cone(X)$. However, every embedded curve in a tree is geodesic. □

The same applies to all Gromov-hyperbolic geodesic metric spaces (e.g., Gromov-hyperbolic groups). Morse lemma fails completely in the case of quasi-geodesics in the Euclidean plane.

7 Boundary extension of QI maps of hyperbolic spaces

Suppose that $X = \mathbb{H}^n$ and $f : X \to X$ is a QI map. Then it sends geodesic rays uniformly close to geodesic rays: $\forall \rho, \exists \rho'$ so that

$$d(f(\rho), \rho') \le D$$

where ρ, ρ' are geodesic rays. Let ξ, ξ' be the limits of the rays ρ, ρ' on the boundary sphere of $\mathbb{H}^n$. Then we set

$$f_\infty(\xi) := \xi'.$$

Exercise 7.1. *The point ξ' depends only on the point ξ and not on the choice of a ray ρ that limits to ξ.*

Thus, we obtain the *boundary extension* of quasi-isometries of $\mathbb{H}^n$ to the boundary sphere S^{n-1}.

Exercise 7.2. $(f \circ g)_\infty = f_\infty \circ g_\infty$ *for all quasi-isometries $f, g : X \to X$.*

Exercise 7.3. *Suppose that $d(f, g) < \infty$, i.e., there exists $C < \infty$ so that*

$$d(f(x), g(x)) \leq C$$

for all $x \in X$. Then $f_\infty = g_\infty$. In particular, if $\bar{f}$ is quasi-inverse to f, then $(\bar{f})_\infty$ is inverse to f_∞.

Our next goal is to see that the extensions f_∞ are continuous, actually, they satisfy some further regularity properties which will be critical for the proof of Tukia's theorem.

Let γ be a geodesic ray in $\mathbb{H}^n$ and $p : \mathbb{H}^n \to \gamma$ be the orthogonal projection (the nearest-point projection). Then for all $x \in \gamma$ (except for the initial point), $H_x := p^{-1}(x)$ is an $n-1$-dimensional hyperbolic subspace of $\mathbb{H}^n$, which is orthogonal to γ.

Lemma 7.4. *Quasi-isometries quasi-commute with the nearest-point projections. More precisely, let $f : \mathbb{H}^n \to \mathbb{H}^n$ be an (L, A)-quasi-isometry. Let γ be a geodesic ray, γ' be a geodesic ray within distance $\leq D(L, A)$ from the quasi-geodesic $f(\gamma)$. Let $p : \mathbb{H}^n \to \gamma, p' : \mathbb{H}^n \to \gamma'$ be nearest-point projections. Then, for some $C = C(L, A)$, we have:*

$$d(fp, p'f) \leq C,$$

i.e.,

$$\forall x \in \mathbb{H}^n, \quad d(fp(x), p'f(x)) \leq C.$$

Let ξ be the limit point of γ. Then, for $x_i \in \gamma$ converging to ξ, the boundary spheres Σ_i of the subspaces $H_{x_i} = p^{-1}(x_i)$, bound round balls $B_i \subset S^{n-1}$ (containing ξ). These balls form a basis of topology at the point $\xi \in S^{n-1}$. Quasi-isometry property of f implies that points $y_i = f(x_i)$ cannot form a bounded sequence in $\mathbb{H}^n$, hence, $\lim y_i = \xi$. Using the above lemma, we see that $f_\infty(B_i)$ are contained in round balls B'_i, whose intersection is the point $\xi' = f_\infty(\xi)$. Thus, f_∞ is continuous and, hence, a homeomorphism. We thus obtain

Lemma 7.5. *For every quasi-isometry $f : \mathbb{H}^n \to \mathbb{H}^n$, the boundary extension f_∞ is a homeomorphism.*

Corollary 7.6. $\mathbb{H}^n$ *is QI to* $\mathbb{H}^m$ *iff* $n = m$.

8 Quasi-actions

The notion of an *action* of a group on a space is replaced, in the context of quasi-isometries, by *quasi-action*. Recall that an *action* of a group G on a set X is a homomorphism $\phi : G \to Aut(X)$, where $Aut(X)$ is the group of bijections $X \to X$. Since quasi-isometries are defined only up to "bounded noise", the concept of a homomorphism has to be modified when we use quasi-isometries.

Definition 8.1. *Let G be a group and X be a metric space. An (L, A) -quasi-action of G on X is a map $\phi : G \to Map(X, X)$, so that:*

- *$\phi(g)$ is an (L, A)-quasi-isometry of X for all $g \in G$.*

- *$d(\phi(1), id_X) \leq A$.*

- *$d(\phi(g_1 g_2), \phi(g_1)\phi(g_2)) \leq A$ for all $g_1, g_2 \in G$.*

Thus, Parts 2 and 3 say that ϕ is "almost" a homomorphism with the error A.

Example 8.2. *Suppose that G is a group and $\phi : G \to \mathbb{R}$ is a function which determines a quasi-action of G on $\mathbb{R}$ by translations ($g \in G$ acts on $\mathbb{R}$ by the translation by $\phi(x)$). Such maps ϕ are called* quasi-morphisms *and they appear frequently in GGT. Many interesting groups do not admit nontrivial homomorphisms of $\mathbb{R}$ but admit unbounded quasimorphisms.*

Here is how quasi-actions appear in the context of QI rigidity problems. Suppose that G_1, G_2 are groups acting isometrically on metric spaces X_1, X_2 and $f : X_1 \to X_2$ is a quasi-isometry with quasi-inverse $\bar{f}$. We then define a quasi-action ϕ of G_2 on X_1 by
$$\phi(g) = \bar{f} \circ g \circ f.$$

Exercise 8.3. *Show that ϕ is indeed a quasi-action.*

For instance, suppose that $X_1 = \mathbb{H}^n$, $\psi : G_1 \curvearrowright X$ is a geometric action, and suppose that G_2 is a group which is QI to G_1 (and, hence, by M-S Lemma, G_2 is QI to X). We then take $X_2 = G_2$ (with a word metric). Then quasi-isometry $f : G_1 \to G_2$ yields a quasi-action $\phi_{f,\psi}$ of G_2 on $\mathbb{H}^n$.

We now apply our extension functor (sending quasi-isometries of $\mathbb{H}^n$ to homeomorphisms of the boundary sphere). Then, Exercises 7.1 amd 7.2 imply:

Corollary 8.4. *Every quasi-action ϕ of a group G on $\mathbb{H}^n$ extends (by $g \mapsto \phi(g)_\infty$) to an action ϕ_∞ of G on S^{n-1} by homeomorphisms.*

Lemma 8.5. *The kernel for the action ϕ_∞ is finite.*

Proof. The kernel of ϕ_∞ consists of the elements $g \in G$ such that $d(\phi(g), id) < \infty$. Since $\phi(g)$ is an (L, A)-quasi-isometry of $\mathbb{H}^n$, it follows from Morse lemma that $d(\phi(g), id) \leq C = C(L, A)$. Thus, such g, as an isometry $G \to G$ moves every point

at most by $C' = C'(L, A)$. However, clearly the set of such elements of G is finite. Hence, $Ker(\phi_\infty)$ is finite as well. $\qquad\square$

Conical limit points of quasi-actions. Suppose that ϕ is a quasi-action of a group G on $\mathbb{H}^n$. A point $\xi \in S^{n-1}$ is called a *conical limit point* for the quasi-action ϕ if the following holds:

For some (equivalently every) geodesic ray $\gamma \subset \mathbb{H}^n$ limiting to ξ, and some (equivalently every) point $x \in \mathbb{H}^n$, there exists a constant $R < \infty$ and a sequence $g_i \in G$ so that:

- $\lim_{i \to \infty} \phi(g)(x) = \xi$.

- $d(\phi(g_i)(x), \gamma) \leq R$ for all i.

In other words, the sequence $\phi(g_i)(x)$ converges to ξ in a closed cone (contained in $\mathbb{H}^n$) with the tip ξ.

Lemma 8.6. *Suppose that $\psi : G \curvearrowright \mathbb{H}^n$ is a geometric action. Then every point of the boundary sphere S^{n-1} is a conical limit point for ψ.*

Lemma 8.7. *Suppose that $G_1 \curvearrowright \mathbb{H}^n$ geometrically, G_2 is a group and $f : \mathbb{H}^n \to G_2$ is a quasi-isometry. Then every point of S^{n-1} is a conical limit point for the corresponding quasi-action $\phi_{f,\psi}$ of G_2 on $\mathbb{H}^n$.*

If ϕ_∞ is a topological action of a group G on S^{n-1} which is obtained by extension of a quasi-action ϕ of G on $\mathbb{H}^n$, then we will say that *conical limit points* of the action $G \curvearrowright S^{n-1}$ are the conical limit points for the quasi-action $G \curvearrowright \mathbb{H}^n$.

9 Quasiconformality of the boundary extension

Can we get a better conclusion than just homeomorphism for the maps f_∞?

I will work in the upper half-space model of $\mathbb{H}^n$. After composing f with isometries of $\mathbb{H}^n$, we can (and will) assume that:

- $\xi \in \mathbb{R}^n$ and γ is the vertical geodesic above 0.

- $\xi' = f_\infty(\xi) \in \mathbb{R}^n$.

- $f_\infty(\infty) = \infty$. In particular, the vertical geodesic γ above ξ maps to a quasi-geodesic within bounded distance from the vertical geodesics γ' above ξ'.

For every $x \in \gamma$, take the subspace $H_x = p^{-1}(x)$ (p is the projection to γ). Then $f(H_x)$ is contained in the slab $S_{y,z} \subset \mathbb{H}^n$ bounded by subspaces H_z, H_y orthogonal to γ', where $d(y, z) \leq C = C(L, A)$. Then the boundary of $S_{y,z}$ in $S^{n-1} = \mathbb{R}^{n-1} \cup \infty$ is a spherical annulus bounded by spheres of radii $R_1 \leq R_2$. We also have:

$$\frac{R_2}{R_1} \leq c = e^C.$$

Thus, the image of the sphere $\Sigma = \partial H_x$ is a "quasi-ellipsoid" of the eccentricity $\leq c$.

I will now change my notation and will use n to denote the dimension of the boundary sphere of the hyperbolic $n + 1$-dimensional space. I will think of S^n as the 1-point compactification of $\mathbb{R}^n$ and will use letters x, y, z, etc., to denote points on $\mathbb{R}^n$. I will also use the notation f for the maps $\mathbb{R}^n \to \mathbb{R}^n$.

Definition 9.1. *A homeomorphism $f : \mathbb{R}^n \to \mathbb{R}^n$ is called c-quasi-symmetric[1], if for all $x, y, z \in \mathbb{R}^n$ so that $0 < |x - y| = |x - z| = r$, we have*

$$\frac{|f(x) - f(y)|}{|f(x) - f(z)|} \leq c. \tag{1}$$

We will think of quasi-symmetric maps as homeomorphisms of $S^n = \mathbb{R}^n \cup \infty$, which send ∞ to itself.

Theorem 9.2. *(P.Tukia, F.Paulin) Every c-quasi-symmetric homeomorphism $S^n \to S^n$ extends to an $(L(c), A(c))$-quasi-isometric map of the hyperbolic space.*

The drawback of this definition is that we are restricted to the maps of $\mathbb{R}^n$ rather than S^n. In particular, we cannot apply this definition to Moebius transformations.

Definition 9.3. *A homeomorphism of S^n is called* quasi-moebius *if it is a composition of a Moebius transformation with a quasi-symmetric map.*

We thus conclude that every (L, A)-quasi-isometry $\mathbb{H}^{n+1} \to \mathbb{H}^{n+1}$ extends to a c-quasi-moebius homeomorphism of the boundary sphere. Unfortunately, this definition of quasi-moebius maps is not particularly useful. One can define instead quasi-moebius maps by requiring that they quasi-preserve the cross-ratio, but then the definition becomes quite cumbersome.

What we will do instead is to take the limit in the inequality (1) as $r \to 0$. Then for every c-quasi-symmetric map f we obtain:

$$\forall x, \quad H_f(x) := \limsup_{r \to 0} \left(\sup_{y,z} \frac{|f(x) - f(y)|}{|f(x) - f(z)|} \right) \leq c. \tag{2}$$

Here, for each $r > 0$ the supremum is taken over y, z so that $r = |x - y| = |x - z|$.

Definition 9.4. *Let U, U' be domains in $\mathbb{R}^n$. Then a homeomorphism $f : U \to U'$ is called quasiconformal if $\sup_{x \in U} H_f(x) < \infty$. A quasiconformal map f is called H-quasiconformal[2] if $H(f) := ess\sup_{x \in U} H_f(x) \leq H$. I will abbreviate quasiconformal to* **qc**.

The number $H(f)$ is called the linear dilatation *of f.*

Thus, every c-quasi-symmetric map is c-quasiconformal. The advantage of quasiconformality is that every Moebius map $f : S^n \to S^n$ is 1-quasiconformal on $S^n \setminus f^{-1}(\infty)$. In particular, all quasi-moebius maps are qc. A more difficult result is:

---

[1]Quasi-symmetric maps can be also defined for general metric spaces, but the definition is more involved.

[2]This definition is slightly different from the standard one, see Appendix 3.

Theorem 9.5. *Every H-quasiconformal map $f : \mathbb{R}^n \to \mathbb{R}^n$ is c-quasi-symmetric for some $c = c(H)$ that depends only on H.*

I will assume from now on that $n \geq 2$ since for $n = 1$ quasiconformality is not quite the right definition.

Analytic properties of qc maps.

1. $H(f \circ g) \leq H(f)H(g)$, $H(f^{-1}) = H(f)$.

2. (J.Väisälä) Every qc map f is differentiable a.e. in $\mathbb{R}^n$. Furthermore, its partial derivatives are in $L_{loc}^n(\mathbb{R}^n)$. In particular, they are measurable functions.

3. (J.Väisälä) Jacobian J_f of qc map f does not vanish a.e. in $\mathbb{R}^n$.

4. Suppose that f is an H-quasiconformal map. For x where $D_x f$ exists and is invertible, we let $\lambda_1 \leq ... \leq \lambda_n$ denote the singular values of the matrix $D_x f$. Then
$$\frac{\lambda_n}{\lambda_1} = H_f(x)$$
Thus, the image of the unit sphere in the tangent space $T_x S^n$ under $D_x f$ is an ellipsoid of eccentricity $\leq H$. This is the geometric interpretation of qc maps: They map infinitesimal spheres to infinitesimal ellipsoids of uniformly bounded eccentricity.

5. **QC Liouville's theorem (F.Gehring and Y.Reshetnyak).** 1-quasiconformal maps are conformal. (Here and in what follows I do not require that conformal maps preserve orientation, only that they preserve angles. Thus, from the viewpoint of complex analysis, we allow holomorphic and antiholomorphic maps of the 2-sphere.)

6. **Convergence property for quasiconformal maps (J.Väisälä).** Let $x, y, z \in S^n$ be three distinct points. A sequence of quasiconformal maps (f_i) is said to be "normalized at $\{x, y, z\}$" if the limits $\lim_i f_i(x), \lim_i f_i(y), \lim f_i(z)$ exist and are all distinct. Then: Every normalized sequence of H-quasiconformal maps contains a subsequence which converges to an H-quasiconformal map.

7. **Semicontinuity of linear dilatation (Tukia; Iwaniec and Martin.** Suppose that (f_i) is a convergent sequence of H-quasiconformal maps so that the sequence of functions H_{f_i} converges to a function H in measure:
$$\forall \epsilon > 0, \lim_{i \to \infty} mes(\{x : |H_{f_i}(x) - H(x)| > \epsilon\}) = 0.$$
(Here mes is the Lebesgue measure on S^n.) Then the sequence (f_i) converges to a qc map f so that $H_f(x) \leq H(x)$ a.e..

10 Quasiconformal groups

Recall that we abbreviate quasiconformal to **qc**.

A group G of homeomorphism of S^n is called (uniformly) quasiconformal if there exists $K < \infty$ so that every $g \in G$ is H-quasiconformal.

Example 10.1. *1. Every conformal (Moebius) group is quasiconformal (take $K = 1$).*

2. Suppose that $f : S^n \to S^n$ is H-quasiconformal and G is a group of conformal transformations of S^n. Then then conjugate group $G_f := fGf^{-1}$ is K^2-quasiconformal.

3. Suppose that ϕ is a quasi-action of a group G on $\mathbb{H}^{n+1}$. Then the extension ϕ_∞ defines an action of G on S^n as a qc group.

4. Conversely, in view of the theorem of Paulin and Tukia, every qc group $G \curvearrowright S^n$ extends to a quasi-action $G \curvearrowright \mathbb{H}^{n+1}$.

D.Sullivan proved that for $n = 2$, every qc group is qc conjugate to a conformal group. This fails for $n \geq 3$. For instance, there are discrete qc groups acting on S^3 which are not isomorphic to any subgroup of Mob_3.

Our goal is to prove

Theorem 10.2. *(P.Tukia, 1986) Suppose that G is a (countable) qc group acting on $S^n, n \geq 2$, so that (almost) every point of S^n is a conical limit point of G. Then G is qc conjugate to a group acting conformally on S^n.*

Once we have this theorem, we obtain:

Theorem 10.3. *Suppose that $G = G_2$ is a group QI to a group G_1 acting geometrically on $\mathbb{H}^{n+1}$ ($n \geq 2$). Then G acts acts geometrically on $\mathbb{H}^{n+1}$.*

Proof. We already know that a quasi-isometry $G_1 \to G_2$ yields a quasi-action ϕ of G on $\mathbb{H}^{n+1}$. Every boundary point of $\mathbb{H}^{n+1}$ is a conical limit point for this quasi-action. We also have a qc extension of the quasi-action ϕ to a qc group action $G \curvearrowright S^n$. Theorem 10.2 yields a qc map h_∞ conjugating the group action $G \curvearrowright S^n$ to a conformal action $\eta : G \curvearrowright S^n$. Every conformal transformation g of S^n extends to a unique isometry $ext(g)$ of $\mathbb{H}^{n+1}$. Thus, we obtain a homomorphism $\rho : G \to Isom(\mathbb{H}^{n+1}), \rho(g) = \text{ext}(\eta(g))$; kernel of ρ has to be finite since the kernel of the action $\phi_\infty : G \curvearrowright S^n$ is finite. We need to verify that the action ρ of G on $\mathbb{H}^{n+1}$ is geometric.

a. **Proper discontinuity.** Suppose that there exists a sequence $g_i \in G$ so that $\lim_i \rho(g_i) = 1$. Then $\lim_i \eta(g_i) = id$ and $\lim_i \phi_\infty(g_i) = id$. Thus, for some (equivalently, every) point $x \in \mathbb{H}^{n+1}$, there exists $C < \infty$ so that

$$d(\phi(g_i)(x), x) \leq C, \forall i.$$

Thus, there exists $C' < \infty$ so that for some (equivalently, every) point $x \in G$, $d(g_i x, x) \leq C'$. Here G acts on itself by left multiplication. However, the set of $g \in G$

so that $d(g_i, 1) \leq C'$, is clearly finite. Thus, the sequence (g_i) consists only of finitely many elements of G and, hence, the action ρ is properly discontinuous.

b. **Cocompactness.** Let $h := ext(h_\infty)$ be an extension of h_∞ to quasi-isometry of $\mathbb{H}^n$. Then

$$g \mapsto \bar{h} \circ \rho(g) \circ h$$

determines a quasi-action ν of G on $\mathbb{H}^{n+1}$ whose extension to S^n is the qc action ϕ_∞. Therefore, there exists a constant C_1 so that for all $g \in G$

$$d(\nu(g), \phi(g)) \leq C_1.$$

Since the action of G on itself was transitive, the quasi-action ϕ of G on $\mathbb{H}^{n+1}$ is cocompact in the sense that there exists a constant C_2 so that for some $x \in \mathbb{H}^{n+1}$,

$$\forall y \in \mathbb{H}^{n+1}, \exists g \in G : d(\phi(g)(x), y) \leq C_2.$$

Since the distance between the quas-actions ϕ and ν is bounded, the quasi-action ν is cocompact too. It follows that the action ρ is cocompact as well. $\qquad\square$

Thus, our objective now is to prove Theorem 10.2

11 Invariant measurable conformal structure for qc groups

Let Γ be group acting conformally on $S^n = \mathbb{R}^n \cup \infty$ and let ds_E^2 be the usual Euclidean metric on $\mathbb{R}^n$. Then conformality of the elements of Γ amounts to saying that for every $g \in \Gamma$, and every $x \in \mathbb{R}^n$ (which does not map to ∞ by g)

$$(D_x g)^T \cdot D_x g$$

is a scalar matrix (scalar multiple of the identity matrix). Here and in what follows, $D_x f$ is the matrix of partial derivatives of f at x. In other words, the product

$$(J_{g,x})^{-\frac{1}{2n}} \cdot (D_x g)^T \cdot D_x g$$

is the identity matrix I. Here $J_{g,x} = det(D_x g)$ is the Jacobian of g at x. This equation describes (in terms of calculus) the fact that the transformation g preserves the conformal structure on S^n.

More generally, suppose that we have a Riemannian metric ds^2 on S^n (given by symmetric positive-definite matrices A_x depending smoothly on $x \in \mathbb{R}^n$). A *conformal structure* on $\mathbb{R}^n$ is the metric ds^2 on $\mathbb{R}^n$ up multiplication by a conformal factor. It is convenient to use *normalized* Riemannian metrics ds^2 on $\mathbb{R}^n$, where we require that $det(A_x) = 1$ for every x. Geometrically speaking, this means that the volume of the unit ball in $T_x(\mathbb{R}^n)$ with respect to the metric ds^2 is the same as the volume ω_n of the unit Euclidean n-ball. Normalization for a general metric A_x is given by

multiplication by $det(A)^{-1/n}$. We then identify space of conformal structures on $\mathbb{R}^n$ with smooth matrix-valued function A_x, where A_x is a positive-definite symmetric matrix with unit determinant.

Suppose that g is a diffeomorphism of S^n. Then the *pull-back* $g^*(ds^2)$ of ds^2 under a diffeomorphism $g : S^n \to S^n$ is given by the symmetric matrices

$$M_x = (D_x g)^T A_{gx} D_x g$$

If A_x was normalized, then, in order to have normalized pull-back $g^\bullet(ds^2)$ we again rescale:

$$B_x := (J_{g,x})^{-\frac{1}{2n}} (D_x g)^T \cdot A_{gx} D_x g.$$

How do we use this in the context of qc maps? Since their partial derivatives are measurable functions on $\mathbb{R}^n$, it makes sense to work with *measurable Riemannian metrics* and *measurable conformal structures* on $\mathbb{R}^n$. (One immediate benefit is that we do not have to worry about the point ∞.) We then work with measurable matrix-valued functions A_x, otherwise, nothing changes. Given a measurable conformal structure μ, we define its *linear dilatation* $H(\mu)$ as the essential supremum of the ratios

$$H(x) := \frac{\sqrt{\lambda_n(x)}}{\sqrt{\lambda_1(x)}},$$

where $\lambda_1(x) \leq ... \leq \lambda_n(x)$ are the eigenvalues of A_x. Geometrically speaking, if $E_x \subset T_x \mathbb{R}^n$ is the unit ball with respect to A_x, then $H(x)$ is the eccentricity of the ellipsoid E_x.

We say that a measurable conformal structure μ is *bounded* if $H(\mu) < \infty$.

We say that a *measurable conformal structure* μ on $\mathbb{R}^n$ is invariant under a qc group G if

$$g^\bullet \mu = \mu, \forall g \in G.$$

In detail:

$$\forall g \in G, \quad (J_{g,x})^{-\frac{1}{2n}} (D_x g)^T \cdot A_{gx} \cdot D_x g = A_x$$

a.e. in $\mathbb{R}^n$.

Theorem 11.1. *(D.Sullivan, P.Tukia) Every qc group acting on $S^n, n \geq 2$, admits a bounded invariant measurable conformal structure.*

Proof. The idea is to start with an arbitrary conformal structure μ_0 on $\mathbb{R}^n$ (say, the Euclidean structure) and then "average" it over $g \in G$. I will prove this only for countable groups G (which is all what we need since we are interested in f.g. groups). Let A_x be the matrix-valued function defining a normalized Riemannian metric on $\mathbb{R}^n$, for instance, we can take $A_x = I$ for all x. Then, since G is countable, for a.e. $x \in \mathbb{R}^n$, we have well-defined matrix-valued function corresponding to $g^*(\mu_0)$ on $T_x \mathbb{R}^n$:

$$A_{g,x} := (J_{g,x})^{-\frac{1}{2n}} (D_x g)^T \cdot A_{gx} \cdot D_x g.$$

For such x we let $E_{g,x}$ denote the unit ball in $T_x\mathbb{R}^n$ with respect to $g^*(\mu_0)$. From the Euclidean viewpoint, $E_{g,x}$ is an ellipsoid of the volume ω_n. This ellipsoid (up to scaling) is the image of the unit ball under the inverse of the derivative $D_x g$. Since g is H-quasiconformal, $E_{g,x}$ has bounded eccentricity, i.e., the ratio of the largest to the smallest axis of this ellipsoid is uniformly bounded independently of x and g. Since volume of $E_{g,x}$ is fixed, it follows that the diameter of the ellipsoid is uniformly bounded above and below.

Let U_x denote the union of the ellipsoids

$$\bigcup_{g \in G} E_{g,x}.$$

This set has diameter $\leq R$ for some R independent of x. Note also that U_x is symmetric (about 0). Note that the family of sets U_x is invariant under the group G:

$$(J_{g,x})^{-1/n} D_x g(U_x) = U_{g(x)}, \quad \forall g \in G.$$

Lemma 11.2. *Given a bounded symmetric subset U of $\mathbb{R}^n$ with nonempty interior, there exists unique ellipsoid $E = E_U$ (centered at 0) of the least volume containing U. The ellipsoid E is called the John-Loewner ellipsoid of U.*

Existence of such an ellipsoid is clear. Uniqueness is not difficult, but not obvious (see Appendix 2). We then let E_x denote the John-Loewner ellipsoid of U_x. This ellipsoid defines a measurable function of x to the space of positive-definite $n \times n$ symmetric matrices. In other words, we obtain a measurable Riemannian metric ν on $\mathbb{R}^n$. Uniqueness of the John-Loewner ellipsoid and G-invariance of the sets U_x implies that the action of G preserves ν_x (up to scaling, of course). One can then get a normalized conformal structure μ by rescaling ν, so that

$$g^\bullet \mu = \mu, \forall g \in G.$$

It remains to show that μ is bounded. Indeed, the length of the major semiaxis of E_x does not exceed R while its volume is $\geq Vol(U_x) \geq \omega_n$. Thus, eccentricity of E_x is uniformly bounded. Hence μ is a bounded measurable conformal structure. $\square$

12 Proof of Tukia's theorem

We are now ready to prove Theorem 10.2. As a warm-up, we consider the easiest case, $n = 2$ (the argument in this case is due to D.Sullivan). In the 2-dimensional case, Theorem 10.2 holds without the conical limit points assumption. Let μ be a bounded measurable conformal structure on S^2 invariant under the group G. *Measurable Riemann mapping theorem* for S^2 states that every bounded measurable conformal structure μ on S^2 is quasiconformally equivalent to the standard conformal structure μ_0 on S^2, i.e., there exists a quasiconformal map $f : S^2 \to S^2$ which sends μ_0 to μ:

$$f^\bullet \mu_0 = \mu.$$

Since a quasiconformal group G preserves μ on S^2, it follows that the group $G_f = fgf^{-1}$ preserves the structure μ_0. Thus, G_f acts as a group of conformal automorphisms of the round sphere, which proves theorem for $n = 2$.

We now consider the case of arbitrary $n \geq 2$.

Definition 12.1. *A function $\eta : \mathbb{R}^n \to \mathbb{R}$ is called* approximately continuous *at a point $x \in \mathbb{R}^n$ if for every $\epsilon > 0$*

$$\lim_{r \to 0} \frac{mes\{y \in B_r(x) : |\eta(x) - \eta(y)| > \epsilon\}}{Vol B_r(x)} = 0.$$

In other words, as we "zoom into" the point x, "most" points $y \in B_r(x)$, have value $\eta(y)$ close to $\eta(x)$, i.e., the rescaled functions $\eta_r(x) := \eta(rx)$ converge in measure to the constant function.

We will need the following result from real analysis:

Lemma 12.2. *(E.Borel) For every L^∞ function η on $\mathbb{R}^n$, a.e. point $x \in \mathbb{R}^n$ is an approximate continuity point of η.*

The functions to which we will apply this lemma are the matrix entries of a (normalized) bounded measurable conformal structure $\mu(x)$ on $\mathbb{R}^n$ (which we will identify with a matrix-valued function A_x). Since μ is bounded and normalized, the matrix entries of $\mu(x)$ will be in L^∞.

We let $\mu(x)$ again denote a bounded normalized measurable conformal structure on $\mathbb{R}^n$ invariant under G. Since a.e. point in $\mathbb{R}^n$ is a conical limit point of G, we will find such point ξ which is also an approximate continuity point for $\mu(x)$.

Then, without loss of generality, we may assume that the point ξ is the origin in $\mathbb{R}^n$ and that $\mu(0) = \mu_0(0)$ is the standard conformal structure on $\mathbb{R}^n$. We will identify $\mathbb{H}^{n+1}$ with the upper half-space $\mathbb{R}^{n+1}_+$. Let $e = e_{n+1} = (0, ..., 0, 1) \in \mathbb{H}^{n+1}$.

To simplify the notation, we let $g(x) = \phi(g)(x)$ denote the quasi-action of the elements $g \in G$ on $\mathbb{H}^{n+1}$.

Since 0 is a conical limit point of G, there exists $C < \infty$ and a sequence $g_i \in G$ so that $\lim_{i \to \infty} g_i(e) = 0$ and
$$d(g_i(e), t_i e) \leq c$$

where d is the hyperbolic metric on $\mathbb{H}^{n+1}$ and $t_i > 0$ is a sequence converging to zero. Let T_i denote the hyperbolic isometry (Euclidean dilation) given by

$$x \mapsto t_i x, x \in \mathbb{H}^{n+1}.$$

Set

$$\tilde{g}_i := T_i^{-1} \circ g_i.$$

Then $d(\tilde{g}_i(e), e) \leq c$ for all i. Furthermore, each $\tilde{g}_i$ is an (L, A)-quasi-isometry of $\mathbb{H}^{n+1}$ for fixed L and A. Thus, the sequence $\tilde{g}_i$ subconverges to a quasiisometry $\tilde{g}$

of $\mathbb{H}^{n+1}$. Accordingly, the sequence of quasiconformal maps f_i which are boundary extensions of $\tilde{g}_i$'s, will subconverge to a quasiconformal mapping f. We also have:

$$\mu_i := f_i^\bullet(\mu) = (T_i^{-1})^\bullet \circ g_i^\bullet(\mu) = (T_i^{-1})^\bullet \mu,$$

since $g^\bullet(\mu) = \mu, \forall g \in G$. Thus,

$$\mu_i(x) = \mu(T_i x) = \mu(t_i x),$$

in other words, the measurable conformal structure μ_i is obtained by "zooming into" the point 0. Since x is an approximate continuity point for μ, the functions $\mu_i(x)$ converge (in measure) to the constant function $\mu_0 = \mu(0)$. Thus, we have the diagram:

$$\begin{array}{ccc} \mu & \xrightarrow{f_i} & \mu_i \\ & & \downarrow \\ \mu & \xrightarrow{f} & \mu_0 \end{array}$$

If we knew that the derivatives Df_i subconverge (in measure) to the derivative of Df, then we would conclude that

$$f^\bullet \mu = \mu_0.$$

Then f would conjugate the group G (preserving μ) to a group G_f preserving μ_0 and, hence, acting conformally on S^n.

However, derivatives of quasiconformal maps (in general), converge only in the distributional sense (actually, even weaker than this), which will not suffice for our purposes. Thus, we have to use a less direct argument below.

We restrict to a certain round ball B in $\mathbb{R}^n$. Since μ is approximately continuous at 0, for every $\epsilon \in (0, \frac{1}{2})$,

$$\|\mu_i(x) - \mu(0)\| < \epsilon$$

away from a subset $W_i \subset B$ of measure $< \epsilon_i$, where $\lim_i \epsilon_i = 0$. Thus, for $x \in W_i$,

$$1 - \epsilon < \lambda_1(x) \leq ... \leq \lambda_n(x) < 1 + \epsilon,$$

where $\lambda_k(x)$ are the eigenvalues of the matrix $A_{i,x}$ of the metric $\mu_i(x)$. Thus,

$$H(\mu_i, x) < \frac{\sqrt{1+\epsilon}}{\sqrt{1-\epsilon}} \leq \sqrt{1 + 4\epsilon} \leq 1 + 2\epsilon.$$

away from subsets W_i. For every $g \in G$, each map $\gamma_i := f_i g f_i^{-1}$ is conformal with respect to the structure μ_i and, hence $(1 + 2\epsilon)$-quasiconformal away from the set W_i. Since $\lim_i mes(W_i) = 0$, we conclude, by the semicontinuity property, that each $\gamma := \lim \gamma_i$ is $(1 + 2\epsilon)$-quasiconformal. Since this holds for arbitrary $\epsilon > 0$ and arbitrary round ball B, we conclude that each γ is is conformal (with respect to the standard conformal structure on S^n).

Thus, the group $\Gamma = fGf^{-1}$ consists of conformal transformations. $\qquad \square$

13   QI rigidity for surface groups

The proof of Tukia's theorem mostly fails for groups QI to the hyperbolic plane. The key reason is that quasi-symmetric maps of the circle are differentiable a.e. but are *not absolutely continuous*. Thus, their derivative could (and, in the interesting cases will) vanish a.e. on the circle.

Nevertheless, the same proof yields: If G is a group QI to the hyperbolic plane, then G acts on S^1 by homeomorphisms with finite kernel K, so that the action is "discrete and cocompact" in the following sense:

Let T denote the set of ordered triples of distinct points on S^1. Thus, T is an open 3-dimensional manifold, one can compute its fundamental group and see that it is infinite cyclic, furthermore, T is homeomorphic to $D^2 \times S^1$. The action $G \curvearrowright S^1$, of course, yields an action $G \curvearrowright T$. Then $G \curvearrowright T$ is *properly discontinuous and cocompact*. The only elements of G that can fix at point in T are the elements of K. Thus, $\Gamma = G/K$ acts freely on T and the quotient T/Γ is a closed 3-dimensional manifold M.

It was proven, in a combination of papers by Tukia, Gabai, Casson and Jungreis in 1988—1994, that such Γ acts geometrically and faithfully on the hyperbolic plane. Their proof was mostly topological. One can now also derive this result from Peremlan's proof of Thurston's geometrization conjecture as follows. The infinite cyclic group $\pi_1(T)$ will be a normal subgroup of $\pi_1(M)$. Then, you look at the list of closed aspherical 3-dimensional manifolds (given by the Geometrization Conjecture) and see that such M has to be a *Seifert manifold*, modelled on one of the geometries $\mathbb{H}^2 \times \mathbb{R}$, $SL(2, \mathbb{R})$, Nil, $\mathbb{E}^3$. In the case of the geometries Nil, $\mathbb{E}^3$, one sees that the quotient of π_1 by normal infinite cyclic subgroup yields a group Γ which is VI to $\mathbb{Z}^2$. Such group cannot act on S^1 so that $\Gamma \curvearrowright T$ is properly discontinuous and cocompact. On the other hand, in the case of the geometries $\mathbb{H}^2 \times \mathbb{R}$, $SL(2, \mathbb{R})$ the quotient by a normal cyclic subgroup will have to be VI to a group acting geometrically on $\mathbb{H}^2$.

14   Appendix 1: Hyperbolic space

The upper half-space model of the hyperbolic n-space $\mathbb{H}^n$ is

$$\mathbb{R}^n_+ = \{(x_1, ... x_n) : x_n > 0\}$$

equipped with the Riemannian metric

$$ds^2 = \frac{|dx|^2}{x_n^2}.$$

Thus, the length of a smooth path $p(t), t \in [0, T]$ in $\mathbb{H}^n$ is given by

$$\int_p ds = \int_0^T \frac{|p'(t)|_e}{p_n(t)} dt.$$

Here $|v|_e$ is the Euclidean norm of a vector v and $p_n(t)$ denotes the n-th coordinate of the point $p(t)$.

The (ideal) *boundary sphere* of $\mathbb{H}^n$ is the sphere $S^{n-1} = \mathbb{R}^{n-1} \cup \infty$, where $\mathbb{R}^{n-1}$ consists of points in $\mathbb{R}^n$ with vanishing last coordinate x_n.

Complete geodesics in $\mathbb{H}^n$ are Euclidean semicircles orthogonal to $\mathbb{R}^{n-1}$ as well as vertical straight lines. For instance, if $p, q \in \mathbb{H}^n$ are points on a common vertical line, then their hyperbolic distance is

$$d(p, q) = |\log(p_n/q_n)|$$

The group of isometries of $\mathbb{H}^n$ is denoted $Isom(\mathbb{H}^n)$. Every isometry of $\mathbb{H}^n$ extends to a *Moebius transformation* of the boundary sphere S^{n-1}. The latter are the conformal diffeomorphisms of S^{n-1} in the sense that they preserve (Euclidean) angles. (I do not assume that conformal transformations preserve orientation.) Conversely, every Moebius transformation of S^{n-1} extends to a unique isometry of $\mathbb{H}^n$.

The group Mob_{n-1} of Moebius transformations of S^{n-1} contains all inversions, all Euclidean isometries of $\mathbb{R}^{n-1}$ and all dilations. (Compositions of Euclidean isometries and dilations are called *similarities*.) In fact, a single inversion together with all similarities of $\mathbb{R}^{n-1}$ generate the full group of Moebius transformations. Furthermore, in every similarity of $\mathbb{R}^{n-1}$ extends to a similarity of $\mathbb{R}^n_+$ in the obvious fashion, so that the extension is an isometry of $\mathbb{H}^n$. Similarly, inversions extend to inversions which are also isometries of $\mathbb{H}^n$.

Exercise 14.1. *Show that the group Mob_{n-1} acts transitively on the set of triples of distinct points in S^{n-1}.*

The key fact of hyperbolic geometry that we will need is that *all triangles in $\mathbb{H}^n$ are δ-thin*, i.e., for every hyperbolic triangle with the sides $\gamma_1, \gamma_2, \gamma_3$, there exists a point $x \in \mathbb{H}^n$ so that

$$d(x, \gamma_i) \leq \delta, i = 1, 2, 3.$$

Here δ is some number ≤ 1.

15 Appendix 2: Least volume ellipsoids

Recall that a closed ellipsoid (with nonempty interior) centered at 0 in $\mathbb{R}^n$ can be described as

$$E = E_A = \{x \in \mathbb{R}^n : \varphi_A(x) = x^T A x \leq 1\}$$

where A is some positive-definite symmetric $n \times n$ matrix. Volume of such ellipsoid is given by the formula

$$Vol(E_A) = \omega_n \left(\det(A)\right)^{-1/2}$$

where ω_n is the volume of the unit ball in $\mathbb{R}^n$. Recall that a subset $X \subset \mathbb{R}^n$ is *centrally-symmetric* if $X = -X$.

Theorem 15.1. *(F. John, 1948) For every compact centrally-symmetric subset $X \subset \mathbb{R}^n$ with nonempty interior, there exists unique ellipsoid $E(X)$ of the least volume containing X. The ellipsoid $E(X)$ is called the* John-Loewner *ellipsoid of X.*

Proof. The existence of $E(X)$ is clear by compactness. We need to prove uniqueness. Consider the function f on the space S_n^+ of positive definite symmetric $n \times n$ matrices, given by

$$f(A) = -\frac{1}{2} \log \det(A).$$

Lemma 15.2. *The function f is strictly convex.*

Proof. Take $A, B \in S_n^+$ and consider the family of matrices $C_t = tA + (1 - t)B$, $0 \leq t \leq 1$. Strict convexity of f is equivalent to strict convexity of f on such line segments of matrices. Since A and B can be simultaneously diagonalised by a matrix M, we obtain:

$$f(D_t) = f(MC_tM^T) = -\log \det(M) - \frac{1}{2}\log \det(C_t) = -\log \det(M) + f(C_t),$$

where D_t is a segment in the space of positive-definite diagonal matrices. Thus, it suffices to prove strict convexity of f on the space of positive-definite diagonal matrices $D = Diag(x_1, ..., x_n)$. Then,

$$f(D) = -\frac{1}{2}\sum_{i=1}^{n} \log(x_i)$$

is strictly convex since log is strictly concave. $\square$

In particular, whenever $V \subset S_n^+$ is a convex subset and $f|V$ is proper, f attains unique minimum on V. Since log is a strictly increasing function, the same uniqueness assertion holds for the function $\det^{-1/2}$ on S_n^+. Let $V = V_X$ denote the set of matrices $C \in S_n^+$ so that $X \subset E_C$. Since $\varphi_A(x)$ is linear as a function of A for any fixed $x \in X$, it follows that V convex. Thus, the least volume ellipsoid containing X is unique. $\square$

16 Appendix 3: Different measures of quasiconformality

Let M be an $n \times n$ invertible matrix with singular values $\lambda_1 \leq ... \leq \lambda_n$. Equivalently, these numbers are the square roots of eigenvalues of the matrix MM^T. The singular value decomposition yields:

$$M = UDiag(\lambda_1, ..., \lambda_n)V$$

where U, V are orthogonal matrices.

We define the following *distortion quantities* for the matrix M:

- Linear dilatation:
$$H(M) := \frac{\lambda_n}{\lambda_1} = \|M\| \cdot \|M^{-1}\|,$$

  where $\|A\|$ is the operator norm of the $n \times n$ matrix A:

  $$\max_{v \in \mathbb{R}^n \setminus 0} \frac{|Av|}{|v|}.$$

- Inner dilatation:
$$H_I(M) := \frac{\lambda_1 .... \lambda_n}{\lambda_n^n} = \frac{|\det(M)|}{\|M\|^n}$$

- Outer dilatation
$$H_O(M) := \frac{\lambda_1^n}{\lambda_1 .... \lambda_n} = \frac{\|M^{-1}\|^{-n}}{|\det(M)|}$$

- Maximal dilatation

$$K(M) := \max(H_I(M), H_O(M)).$$

Exercise 16.1.
$$(H(M))^{n/2} \leq K(M) \leq (H(M))^{n-1}$$

Hint: It suffices to consider the case when $M = Diag(\lambda_1, ... \lambda_n)$ is a diagonal matrix.

As we saw, qc homeomorphisms are the ones which send infinitesimal spheres to infinitesimal ellipsoids of uniformly bounded eccentricity. The usual measure of quasiconformality of a qc map f is its *maximal distortion* (or *maximal dilatation*) $K(f)$, defined as
$$K(f) := ess\sup_x K(D_x(f))$$

where the essential supremum is taken over all x in the domain of f. Here $D_x f$ is the derivative of f at x (Jacobian matrix). See e.g. J.Väisälä's book [4]. A map f is called K-quasiconformal if $K(f) \leq K$.

In contrast, the definition of quasiconformality used in these lectures is: f is H-quasiconformal if
$$H(f) := ess\sup_x H(D_x f) \leq H.$$

To relate to the two definition we observe that

$$1 \leq (H(f))^{n/2} \leq K(f) \leq (H(f))^{n-1}.$$

In particular, $K(f) = 1$ if and only if $H(f) = 1$.

References

[1] T. Iwaniec, G. Martin, "Geometric Function Theory and Non-linear Analysis", Oxford Univ. Press, 2001.

[2] D. Sullivan, *On the ergodic theory at infinity of an arbitrary discrete group of hyperbolic motions*, in "Riemann Surfaces and Related Topics": Proceedings of the 1978 Stony Brook Conference, I. Kra and B. Maskit, eds., Ann. Math. Studies 97, Princeton University Press, 1981, pp. 465–496.

[3] P. Tukia, *On quasiconformal groups*, J. d'Anal. Math, vol 46 (1986), pp. 318–346.

[4] J.Väisälä, "Lectures on n-dimensional quasiconformal mappings", Springer Lecture Notes in Math, vol 229, 1961.