

RESEARCH ARTICLE

Stability for hyperbolic groups acting on boundary spheres

Kathryn Mann¹ and Jason Fox Manning²

¹Department of Mathematics, Cornell University, Ithaca, NY 14853, USA; E-mail: k.mann@cornell.edu.

²Department of Mathematics, Cornell University, Ithaca, NY 14853, USA; E-mail: jfmanning@cornell.edu.

Received: 27 July 2023; Revised: 31 July 2023; Accepted: 13 August 2023

2020 Mathematics Subject Classification: Primary – 20F67; Secondary – 37B05, 57S25

Abstract

A hyperbolic group G acts by homeomorphisms on its Gromov boundary. We show that if ∂G is a topological n -sphere, the action is *topologically stable* in the dynamical sense: any nearby action is semi-conjugate to the standard boundary action.

1. Introduction

Calabi–Weil local rigidity [Cal61, Wei62] (an important precursor to Mostow [Mos73] rigidity) states that, for $n \geq 3$, the action of the fundamental group of a hyperbolic n -manifold by conformal maps on the boundary sphere S^{n-1} is locally rigid: any nearby conformal action is conjugate in $SO^+(n, 1)$ to the original action. Inspired by this, we investigate rigidity for the actions on boundary spheres of the broader class of all Gromov hyperbolic groups with sphere boundary. These boundaries do not typically admit a natural conformal or even a C^1 structure, so the relevant notion of local stability is that from topological dynamics. Recall that an action $\rho_0: G \rightarrow \text{Homeo}(X)$ of a group G on a topological space X is a *topological factor* of an action $\rho: G \rightarrow \text{Homeo}(Y)$ if there is a surjective, continuous map $h: Y \rightarrow X$, such that $h \circ \rho = \rho_0 \circ h$. Such a map h is called a *semiconjugacy*. An action of a group G on a topological space X is *topologically stable* or C^0 *stable* if it is a factor of any sufficiently close action in the compact-open topology on $\text{Hom}(G, \text{Homeo}(X))$. We prove the following.

Theorem 1.1 (Topological stability). *Let G be a hyperbolic group with sphere boundary. Then the action of G on ∂G is topologically stable. More precisely, given any neighborhood V of the identity in the space of continuous self-maps of S^n , there exists a neighborhood U of the standard boundary action in $\text{Hom}(G, \text{Homeo}(S^n))$, such that any representation in U has ρ_0 as a factor, with semiconjugacy contained in V .*

In parallel with Calabi–Weil rigidity, this says that these boundary actions exhibit the strongest possible form of local rigidity. While there is overlap in the groups considered (fundamental groups of closed hyperbolic manifolds are Gromov hyperbolic), our result is neither a special case nor a generalization of the classical case. We consider a much broader space of deformations – actions by homeomorphisms rather than conformal maps – but semiconjugacy is of course weaker than conformal conjugacy, which one cannot hope for when considering general continuous deformations (see, in particular, the examples of [BM19, Section 4]).

History and related results

“Stability from hyperbolicity” is an important and recurring theme in dynamical systems. For instance, hyperbolic (Anosov) diffeomorphisms are topologically stable, thanks to the well known shadowing lemma. However, in much of the existing literature, hyperbolicity is described using some smooth or at least C^1 structure, while the actions we consider are typically not differentiable, only having Hölder regularity.

Regarding boundary actions of groups, Sullivan’s 1985 *Structural stability implies hyperbolicity for Kleinian groups* [Sul85], characterizes convex-cocompact subgroups of $\mathrm{PSL}(2, \mathbb{C})$ as those subgroups whose action on their limit set is stable under C^1 perturbations. Sullivan uses the fact that group elements expand neighborhoods of points to produce a coding of orbits that is insensitive to perturbation. This technique was recently generalized by Kapovich–Kim–Lee [KKL] to a much broader setting, including Lipschitz perturbations of many group actions on metric spaces which satisfy a generalized version of Sullivan’s expansivity condition.

Matsumoto [Mat87] gives a more robust form of rigidity for the actions of fundamental groups of compact surfaces on their boundary at infinity. In this case, the boundary is a topological circle, and Matsumoto’s work implies that any deformation of such a boundary action is semiconjugate to the original action. Motivated by this, Bowden and the first author [BM19] studied the actions of the fundamental groups of compact Riemannian manifolds on their boundaries at infinity, showing these satisfy a form of local rigidity. Again hyperbolicity played a role, this time in the form of the Anosov property of geodesic flow on such negatively curved manifolds. Theorem 1.1 generalizes aspects of both Sullivan’s and Matsumoto’s program. While hyperbolic groups acting on their boundaries are among the examples studied by Kapovich–Kim–Lee, their methods only apply to perturbations which continue to have Sullivan’s expansivity, for instance, Lipschitz-close actions. General C^0 perturbations need not be Lipschitz close, so Sullivan’s coding no longer applies, and we need an entirely new method of proof. Our strategy is more in the spirit of [BM19] but uses large-scale geometry in place of the Riemannian manifold structure and Anosov geodesic flow.

Our focus on spheres is motivated in part by the fact that these are the most homogeneous group boundaries. At the other end of the spectrum, Kapovich and Kleiner [KK00] constructed hyperbolic groups that are *boundary rigid*, in the sense that any homeomorphism of the boundary comes from the action of an element of the group. These groups trivially satisfy local rigidity, since $\mathrm{Homeo}(\partial G) \cong G$ is discrete. By contrast, homeomorphisms of the sphere are very easy to perturb, each having an infinite dimensional family of deformations. The reader may consult [CP93] or [KB02] for more background on the dynamics of hyperbolic groups acting on their boundaries.

Scope

Bartels, Lück and Weinberger [BLW10, Example 5.2] give, for all $k \geq 2$, examples of torsion-free hyperbolic groups G with $\partial G = S^{4k-1}$ that are not the fundamental group of any smooth, closed, aspherical manifold (note that such an example with $\partial G = S^2$ would give a counterexample to the Cannon Conjecture). These examples show that, even in the torsion free case, Theorem 1.1 is a strict generalization of the work of [BM19] on Riemannian manifold fundamental groups. Of course, groups with torsion provide numerous other examples, and the tools introduced within the large scale geometric framework of the proof should be of independent interest.

Outline

The broad strategy of the proof is to translate the data of a G -action on S^n into a G -action on a sphere bundle over a particular space quasi-isometric to G , then show that nearby actions can be related by a G -equivariant map between their respective bundles that is close to the identity on large compact sets. This lets us promote metrically stable notions in coarse negative curvature (such as the property of a subset being bounded Hausdorff distance from a geodesic) into stability for the group action.

In Section 2, we collect general results and preliminary lemmas on hyperbolic metric spaces. In Section 3, we construct the bundles and equivariant map advertised above, in the broader context of (not necessarily hyperbolic) groups acting on manifolds, which is the natural setting for this technique.

To apply this technique to the proof of Theorem 1.1, we need to find a suitably nice space X with a proper, free, cocompact action of G by isometries. If G is torsion free, the space of distinct triples in ∂G is a natural choice, but if G has torsion, the action of G on triples may not be free. We remedy this in Sections 4 and 5, first reducing to the case where G acts faithfully on ∂G , and then showing that one may remove a small neighborhood of the space of fixed boundary triples without losing too much geometry, giving a suitable space to use in the rest of the proof. We also prove several technical lemmas on the triple space for reference in later sections.

Section 6 sets up the main proof and specifies a neighborhood of the boundary action where Theorem 1.1 holds. The bundles from Section 3 come with natural topological foliations, and Section 6.2 shows that the image of one of these foliations in the source space (whose leaves are parameterized by points of ∂G) intersects leaves in the target along coarsely geodesic sets. Section 7 shows that the endpoints of these coarse geodesics depend only on the original leaf, thus giving a map h from the leaf space of one foliation to the Gromov boundary of X . Both the leaf space and the Gromov boundary of X are canonically homeomorphic to ∂G , and the map h is our semiconjugacy.

2. Background

We set notation, collect some general results on hyperbolic metric spaces, and prove some preliminary lemmas needed for the main theorem.

2.1. Setup

We fix the following notation. G denotes a nonelementary hyperbolic group; in this section, we do not require ∂G to be a sphere. We fix a generating set \mathcal{S} , which gives us a Cayley graph Γ and metric d_Γ on Γ . Vertices of Γ are identified with group elements. In particular, the identity e is a vertex of Γ . The metric d_Γ is ν -hyperbolic (in the sense that geodesic triangles are ν -thin) for some $\nu > 0$. We will fix a constant $\delta \geq \nu$ with some other convenient properties later. The Gromov boundary of the group is denoted ∂G , this is of course equal to the Gromov boundary of Γ .

We write $(x | y)_z$ for the Gromov product of x and y at z . The point z must lie in Γ , but x, y may be in $\Gamma \cup \partial G$, using the standard definition of the Gromov product at infinity (see e.g. [BH99, III.H.3.15]), as follows:

$$(x | y)_p = \sup \left\{ \liminf_{i,j \rightarrow \infty} (x_i | y_j)_p \mid \lim_{i \rightarrow \infty} x_i = x, \lim_{i \rightarrow \infty} y_i = y \right\}.$$

We also fix a *visual metric* d_{vis} on ∂G . This means a metric so that there are constants $\lambda > 1$ and $k_2 > k_1 > 0$ satisfying, for all $a, b \in \partial G$,

$$k_1 \lambda^{-(a | b)_e} \leq d_{\text{vis}}(a, b) \leq k_2 \lambda^{-(a | b)_e} \tag{1}$$

(see [BH99, III.H.3] or [GdlH90, 7.3] for more details, including the existence of such a metric). Unless otherwise specified, all metric notions in ∂G (such as balls $B_r(p)$) will be defined using this visual metric. Occasionally, when specializing to $\partial G = S^n$, we also make use of the standard round metric on S^n .

We will need to use the following lemma for estimating the Gromov product of points at infinity in a ν -hyperbolic space.

Lemma 2.1. *Let M be ν -hyperbolic, and let $p \in M$. Let $\alpha, \beta \in \partial M$ be represented by geodesic rays $\gamma_\alpha, \gamma_\beta$ starting at p . For any points $a \in \gamma_\alpha, b \in \gamma_\beta$,*

$$(\alpha | \beta)_p \geq (a | b)_p - 2\nu.$$

Proof. From the definition of Gromov product at infinity, it follows that

$$(\alpha | \beta)_p \geq \liminf_{s,t \rightarrow \infty} (\gamma_\alpha(s) | \gamma_\beta(t))_p.$$

To simplify notation, consider p as a base point and write $|\cdot|$ for $d_M(p, \cdot)$. Suppose that c is a point on γ_α and d a point on γ_β so that $|c| > |a|$ and $|d| > |b|$. We want to show that $(c | d)_p \geq (a | b)_p - 2\nu$.

Consider the triple of Gromov products $(c | d)_p$, $(b | d)_p = |b|$, and $(c | b)_p \leq |b|$. As is well known, it follows from ν -hyperbolicity that any one of such a triple is bounded below by ν less than the minimum of the other two, so we have

$$(c | d)_p \geq (c | b)_p - \nu. \tag{2}$$

Considering next the triple $(c | b)_p$, $(c | a)_p = |a|$, and $(a | b)_p \leq |a|$, we conclude

$$(c | b)_p \geq (a | b)_p - \nu. \tag{3}$$

Together (2) and (3) give $(c | d)_p \geq (a | b)_p - 2\nu$, as desired. □

2.2. The space of triples

Write Ξ for the space of ordered distinct triples of points in the Gromov boundary $\partial\Gamma = \partial G$. We use the following well known property.

Proposition 2.2 (Convergence group property (see [Bow98, Tuk98] and [Gro87] Section 8.2.M)). *G acts properly discontinuously and cocompactly on Ξ , and each point $a \in \partial G$ is a conical limit point, meaning that there exists $\{g_i\}_{i \in \mathbb{N}} \subset G$ and $p \neq q \in \partial G$, such that $g_i(a) \rightarrow p$ and $g_i(z) \rightarrow q$ for all $z \in \partial G - \{a\}$.*

The following definition can be thought of as giving a coarse projection map from Ξ to Γ (compare [Gro87, Section 8.2.K]). When S is a subset of a metric space, we use the notation $N_r(S)$ to indicate the open r -neighborhood of S .

Definition 2.3 (Coarse projection). For each $r > 0$, we define a projection map π_r from Ξ to subgraphs of Γ as follows. For $(a, b, c) \in \Xi$, let $\mathcal{G}(a, b, c)$ be the set of geodesics in Γ with endpoints in $\{a, b, c\}$. For each $r > 0$, define $\pi_r(a, b, c) \subset \Gamma$ to be the smallest subgraph of Γ containing $\bigcap_{\gamma \in \mathcal{G}(a,b,c)} N_{r-1}(\gamma)$.

If Z is a subset of Ξ , we define $\pi_r(Z) = \bigcup_{z \in Z} \pi_r(z)$. For $s \in \Gamma$, we define $\pi_r^{-1}(s) = \{(a, b, c) \mid s \in \pi_r(a, b, c)\}$; for $S \subset \Gamma$, define $\pi_r^{-1}(S) = \bigcup_{s \in S} \pi_r^{-1}(s)$.

Remark 2.4. If r is sufficiently large (depending on the hyperbolicity constant of Γ), then $\pi_r(a, b, c)$ is always nonempty. Moreover, for any $x \in \pi_r(a, b, c)$ and any geodesic γ with endpoints in $\{a, b, c\}$, we have $d_\Gamma(x, \gamma) \leq r$. We will make frequent use of this estimate.

Lemma 2.5. *For every $r \geq 0$, there is a $Q(r) \geq 0$ so $\text{diam}(\pi_r(a, b, c)) \leq Q(r)$ for all $(a, b, c) \in \Xi$.*

Proof. Recall Γ is ν -hyperbolic, meaning that triangles are ν -thin. Let $(a, b, c) \in \Xi$ be given. Fix biinfinite geodesics $[a, b]$, $[b, c]$, and $[a, c]$ in Γ . Approximate this ideal triangle by a triangle in Γ by choosing points a', b' and $c' \in \Gamma$ on the geodesics $[a, b]$, $[b, c]$, and $[a, c]$, respectively, satisfying $d_\Gamma(a', [a, c]) < \nu$ and $d_\Gamma(a', [b, c]) > r + 2\nu$, and such that the same two inequalities also hold when the letters a, b, c are cyclically permuted. Then $\pi_r(a, b, c)$ is a subset of

$$S_r := N_{r+\nu}([a', b']) \cap N_{r+\nu}([b', c']) \cap N_{r+\nu}([c', a'])$$

so it suffices to show this set has diameter bounded by $Q(r)$, for some suitable function Q .

Let p denote the map from the triangle with sides $[a', b']$, $[b', c']$, and $[a', c']$ to a tripod witnessing that the triangle is ν -thin, and let Z be the preimage of the center, this is a set of three points with

diameter at most ν . We now claim that S_r lies in the $3r + 5\nu$ -neighborhood of Z , which is enough to prove the lemma. To prove the claim, suppose $s \in S_r$, so there exist points x_1, x_2 , and x_3 on $[a', b']$, $[b', c']$, and $[a', c']$, respectively, with $d_\Gamma(x_i, s) < r + \nu$. Then for any $i = 1, 2, 3$, there exists some j so that $p(x_i)$ and $p(x_j)$ lie on different prongs of the tripod, so there is a path between them passing through the midpoint m of the tripod. Thus, we have

$$d_\Gamma(x_i, Z) = d_\Gamma(p(x_i), m) \leq d_\Gamma(p(x_i), p(x_j)) \leq d_\Gamma(x_i, x_j) + 2\nu \leq 2r + 4\nu,$$

where the last inequality follows from the fact that $d_\Gamma(x_i, s) < r + \nu$. This proves the claim. □

Lemma 2.6. *Let $r \geq 0$. For any compact $K \subset \Xi$, the set $\pi_r(K)$ is bounded.*

Proof. Let $K \subset \Xi$ be compact. Increasing r makes $\pi_r(K)$ larger. Using Remark 2.4, we can therefore assume that for every (a, b, c) , there is some $x \in \pi_r(a, b, c)$.

For each $(a, b, c) \in K$, there is an open neighborhood U of (a, b, c) in Ξ , such that, for each point $(u, v, w) \in U$, any geodesics joining points in $\{u, v, w\}$ come within $2\nu + r$ of the point x . In particular, $x \in \pi_{r+2\nu}(u, v, w)$. If $y \in \pi_r(u, v, w)$, then $d_\Gamma(x, y) \leq Q(r + 2\nu)$, so $\pi_r(U)$ has diameter at most $2Q(r + 2\nu)$.

By compactness, we can cover K with finitely many neighborhoods U as in the last paragraph, so $\pi_r(K)$ is bounded. □

It will be convenient to choose a hyperbolicity constant for Γ that simultaneously satisfies several properties. The properties we use are collected in the following lemma.

Lemma 2.7. *There exists $\delta > 0$ so that all of the following hold:*

- ($\delta 1$) *Every geodesic triangle in Γ is δ -thin.*
- ($\delta 2$) *Every geodesic bigon or triangle with vertices in $\Gamma \cup \partial G$ is δ -slim.*
- ($\delta 3$) *For any point $p \in \Gamma$, and any $a, b, c \in \Gamma \cup \partial G$,*

$$(a | b)_p \geq \min\{(a | c)_p, (b | c)_p\} - \delta.$$

($\delta 4$) *For all $p \in \Gamma$, $\pi_\delta^{-1}(p)$ is nonempty.*

($\delta 5$) *The set $\pi_\delta(\{a\} \times \{b\} \times (\partial G - \{a, b\}))$ contains every geodesic joining a to b .*

Proof. Since Γ is ν -hyperbolic, items ($\delta 1$) and ($\delta 2$) hold for any $\delta \geq 2\nu$. For item ($\delta 3$), see [BH99, III.H.3.17.(4)]. Item ($\delta 4$) follows from G -equivariance and the fact that $\pi_\delta(a, b, c)$ is nonempty when δ is large enough. For ($\delta 5$), suppose we are given a point z on a geodesic γ joining a and b , take $c \in \partial G$ minimizing $\max\{(a | c)_z, (b | c)_z\}$. Cocompactness of the action of G allows one to bound this minimum from above, independently of a, b , and c , and this can be used to give an upper bound on the distance from z to any geodesic joining a or b with c . □

Notation 2.8. For the rest of the paper, we fix some $\delta > 0$ so the conclusions of Lemma 2.7 hold, and denote the coarse projection π_δ by π .

Definition 2.9 (Minimum separation). For $x = (a, b, c) \in \Xi$, we define

$$\text{minsep}(x) = \min\{d_{\text{vis}}(a, b), d_{\text{vis}}(a, c), d_{\text{vis}}(b, c)\}.$$

Notice that $1/\text{minsep}$ is a proper function on X , so minsep is bounded away from zero on any compact set. For a subset $D \subset \Xi$, we define $\text{minsep}(D) = \inf\{\text{minsep}(x) \mid x \in D\}$.

2.3. A criterion for a set to be close to a geodesic

The following lemma gives a criterion for a piecewise geodesic curve to be close to a geodesic. There are various similar statements in the literature (e.g. [Min05, Lemma 4.2], [BH99, Section III.H.1.13]),

but this form will be convenient for us. We use it to prove Lemma 2.12 which is the main technical ingredient of this section.

Lemma 2.10. *Let X be a δ -hyperbolic geodesic metric space, and let $l > 0$. Suppose that c is a piecewise geodesic in X made of segments of length greater than $2l + 8\delta$, with Gromov products in the corners at most l . Let γ be a geodesic with the same endpoints as c . The Hausdorff distance between γ and c is at most $l + 4\delta$.*

We remark that this lemma only uses δ -hyperbolicity, and not the other properties from Lemma 2.7.

Proof. A standard argument, using only the fact that γ is a geodesic in a δ -hyperbolic space, shows that it is enough to prove c is contained in the closed $(l + 3\delta)$ -neighborhood of γ . We will not give the details, as this is classical. We write c as a concatenation $c_1 \cdots c_k$ of geodesics, so each c_i joins some p_{i-1} to some p_i . The endpoints of γ are p_0 and p_k . If $k \leq 2$, we are done by slimness of triangles, so we assume $k \geq 3$.

Let x be the farthest point from γ on c , and let $M = d(x, \gamma)$. Without loss of generality, we suppose that $M > 2\delta$. It is then straightforward to show that x is within 2δ of some breakpoint p_i (consider the triangle made up of the segment c_i containing x , together with geodesics joining the endpoints of c_i to a point $x' \in \gamma$ closest to x).

Since $M > 2\delta$, the breakpoint p_i cannot be either endpoint of the geodesic γ ; in particular, $i \notin \{0, k\}$. There are two cases, depending on whether or not $i \in \{1, k-1\}$.

We suppose $i \notin \{1, k-1\}$; the case $i \in \{1, k-1\}$ is similar but easier. By the assumption that segments are long, $d(x, \{p_{i\pm 1}\}) > 2l + 6\delta$. Choose a geodesic σ joining p_{i-1} to p_{i+1} . By the assumption on Gromov products in the corners, we have $(p_{i-1} | p_{i+1})_{p_i} \leq l$. It follows that $d(x, \sigma) \leq l + \delta$. Let y be a closest point to p_{i-1} in γ , and let z be a closest point to p_{i+1} in γ . Choose geodesics $[y, z] \subset \gamma$, $[p_{i-1}, y]$, and $[p_{i+1}, z]$. The point x lies within $l + 3\delta$ of some point w on the union of these three geodesics. We claim that $w \in [y, z]$, so we have $M \leq l + 3\delta$.

Indeed, suppose that $w \in [p_{i-1}, y]$ (the case $w \in [p_{i+1}, z]$ being identical). Now we have

$$\begin{aligned} 0 \leq d(x, y) - d(p_{i-1}, y) &\leq d(x, w) + d(w, y) - (d(p_{i-1}, w) + d(w, y)) \\ &= d(x, w) - d(p_{i-1}, w) \\ &\leq d(x, w) - (d(x, p_{i-1}) - d(x, w)) \\ &= 2d(x, w) - d(x, p_{i-1}) \\ &\leq 2(l + 3\delta) - d(x, p_{i-1}) < 0 \end{aligned}$$

a contradiction. We have thus established that $M \leq l + 3\delta$, and so c lies in the $l + 3\delta$ -neighborhood of γ . \square

Definition 2.11. Let $r > 0$, and let M be a metric space. A subset S of M is r -connected if any two points p, q of S can be connected by a chain of points in S ,

$$p = p_0, p_1, \dots, p_k = q,$$

so that $d_M(p_i, p_{i+1}) \leq r$ for all i . An r -connected component of S is a maximal subset of S which is r -connected.

Lemma 2.12. *Let $H > 0$, and let $R > 24H + 16\delta$. Let $S \subset \Gamma$ be a $\frac{R}{4}$ -connected set so that for every $s \in S$, there is a biinfinite geodesic γ_s satisfying:*

$$S \cap B_R(s) \subset N_H(\gamma_s), \text{ and } \gamma_s \cap B_R(s) \subset N_H(S).$$

Then there is an oriented biinfinite geodesic γ so that

1. $d_{\text{Haus}}(\gamma, S) \leq 3H + 6\delta$; and
2. for every $s \in S$, we may orient γ_s so the Gromov products $(\gamma^{+\infty} | \gamma_s^{+\infty})_s$ and $(\gamma^{-\infty} | \gamma_s^{-\infty})_s$ are bounded below by $R - (4H + 10\delta)$.

Proof. Choose any s_0 in S , and let $\gamma_0 = \gamma_{s_0}$ be a biinfinite geodesic as in the hypothesis, parameterized so that $\gamma_0(0)$ is within H of s_0 . Since the points $\gamma_0(\pm \frac{R}{2})$ lie in $B_R(s)$, there are points $s_{\pm 1} \in S$ whose distances from $\gamma_0(\pm \frac{R}{2})$ are at most H . Since $d_\Gamma(s_0, s_{\pm 1}) \leq \frac{R}{2} + 2H$ and $d_\Gamma(s_{-1}, s_1) \geq R - 2H$, we deduce $(s_{-1} | s_1)_{s_0} \leq 3H$. In particular, we have the following estimates:

$$d_\Gamma(s_0, s_{\pm 1}) \geq \frac{R}{2} - 2H \text{ and } (s_{-1} | s_1)_{s_0} \leq 3H. \tag{4}$$

Now, we inductively find s_i and γ_i for all integers i .

For clarity, we focus on $i > 1$. The construction for $i < 0$ is entirely analogous. Suppose we have chosen points $s_{-1}, s_0, \dots, s_{i-1}$ in S , and that for each positive $j \leq i - 1$, we have chosen a biinfinite geodesic $\gamma_j = \gamma_{s_j}$, and some t_j within $4H$ of $-\frac{R}{2}$ so that

$$d_\Gamma(\gamma_{i-1}(0), s_{i-1}) \leq H, \text{ and} \\ d_\Gamma(\gamma_{i-1}(t_{i-1}), s_{i-2}) \leq H.$$

Since $R > 8H$, the number t_{i-1} is negative. The point $\gamma_{i-1}(\frac{R}{2})$ lies in the R -ball around s_{i-1} , so we may choose a point $s_i \in S$ so that $d_\Gamma(s_i, \gamma_{i-1}(\frac{R}{2})) \leq H$. Let γ_i be the geodesic γ_{s_i} provided by the hypothesis of the lemma. We can assume that $\gamma_i(0)$ is within H of s_i . The distance $d_\Gamma(s_{i-1}, s_i)$ differs from $\frac{R}{2}$ by at most $2H$. Thus, for some t_i of absolute value in $[\frac{R}{2} - 4H, \frac{R}{2} + 4H]$, we have $d_\Gamma(\gamma_i(t_i), s_{i-1}) \leq H$. We parameterize γ_i so that $t_i < 0$. This completes the inductive construction.

From the construction, we have

$$d_\Gamma(s_i, s_{i+1}) \leq \frac{R}{2} + 2H \tag{5}$$

and

$$d_\Gamma(s_{i-1}, s_{i+1}) \geq \frac{R}{2} + |t_i| - 2H \geq R - 6H$$

(the lower bound when $i = 0$ is slightly better). This implies a bound on Gromov products

$$(s_{i-1} | s_{i+1})_{s_i} \leq 5H. \tag{6}$$

Let c_k be a piecewise geodesic formed by concatenating geodesics

$$[s_{-k}, s_{-k+1}] \cdots [s_{k-1}, s_k].$$

We verify the hypotheses of Lemma 2.10 with $l = 5H$. The inequality (6) gives the bound on Gromov products in the corners. The inequality (5) gives that the segments $[s_i, s_{i+1}]$ have length at least $\frac{R}{2} - 2H > 10H + 8\delta = 2l + 8\delta$ as required. Thus, if β_k is the geodesic joining the endpoints of c_k , we have $d_{\text{Haus}}(c_k, \beta_k) \leq H + 4\delta$.

Since Γ is proper, and the geodesics β_k all pass through the $(H + 4\delta)$ -ball about s_0 , they subconverge to a biinfinite geodesic γ . Notice that all the segments $[s_i, s_{i+1}]$ lie in the $(H + 4\delta)$ -neighborhood of γ . We will show this γ satisfies the conclusions of the lemma.

If $s \in S$, then there is a $\frac{R}{4}$ -coarse path joining s_0 to s , that is to say there exist points $s_0 = p_0, p_1, \dots, p_k = s$ in S satisfying

$$d_\Gamma(p_i, p_{i+1}) \leq \frac{R}{4}, \forall i.$$

We first claim that for each p_i , there is some $s_{j(i)}$ with $d_\Gamma(s_{j(i)}, p_i) \leq \frac{R}{4} + 2H$. Clearly this is true for p_0 . Arguing by induction, we see that $d_\Gamma(p_{i+1}, s_{j(i)})$ is at most $\frac{R}{2} + 2H$. In particular, there is a point q on $\gamma_{s_{j(i)}}$ within H of p_{i+1} . Recalling that $\gamma_{s_{j(i)}}(0)$ lies within H of $s_{j(i)}$, we see that $q = \gamma_{s_{j(i)}}(t)$ for some t with

$$|t| \leq \frac{R}{2} + 3H < \frac{3}{4}R.$$

Thus, for some $t' \in \{-\frac{R}{2}, 0, \frac{R}{2}\}$, we have $|t - t'| \leq \frac{R}{4}$. Since the points $\gamma_{s_{j(i)}}(\pm\frac{R}{2})$ are within H of $s_{j(i)\pm 1}$, there is some $s_{j(i+1)} \in \{s_{j(i)-1}, s_{j(i)}, s_{j(i)+1}\}$ so that $d_\Gamma(p_{i+1}, s_{j(i+1)}) \leq 2H + \frac{R}{4}$, as desired.

Now let $s_j = s_{j(k)}$, so we have $d_\Gamma(s, s_j) \leq \frac{R}{4} + 2H$, and let q be a closest point to s on γ_j . Without loss of generality, we suppose that $q = \gamma_j(t)$ for $t \geq 0$. Any quadrilateral with corners $s_j, s_{j+1}, \gamma_j(0), \gamma_j(\frac{R}{2})$ is 2δ -thin, so there is a point r on $[s_j, s_{j+1}]$ within $H + 2\delta$ of q . This point r is within $H + 4\delta$ of some point z on γ . Adding up the constants, we have

$$\begin{aligned} d_\Gamma(s, z) &\leq d_\Gamma(s, q) + d_\Gamma(q, r) + d_\Gamma(r, z) \\ &\leq H + H + 2\delta + H + 4\delta = 3H + 6\delta. \end{aligned}$$

This shows

$$S \subset N_{3H+6\delta}(\gamma). \tag{7}$$

Conversely, let $x \in \gamma$. Then $x \in \beta_k$ for some k , and so for some i , there is a point $y \in [s_i, s_{i+1}]$ with $d_\Gamma(x, y) \leq H + 4\delta$. This point is within $H + 2\delta$ of a point on γ_i , which is within H of a point of S , so we have

$$\gamma \subset N_{3H+6\delta}(S). \tag{8}$$

Together, (7) and (8) imply the first statement of the lemma, that is to say the bound on Hausdorff distance. It remains to show the statement about Gromov products. Breaking symmetry, we consider just the ray $\gamma_s|[0, \infty)$. Let y' be a point on $\gamma_s|[0, \infty)$ at distance R from s . Let $s' \in S$ be a point within H of y' , and let z' be a point on γ within $3H + 6\delta$ of s' . Let α be a ray starting at s with limit point $\gamma_s^{+\infty}$, and let β be a ray starting at s with limit point $\gamma^{+\infty}$. There are points y on α and z on β which are within δ of y', z' , respectively. We have $d_\Gamma(s, y) \geq R - \delta$, $d_\Gamma(s, z) \geq R - (4H + 7\delta)$ and $d_\Gamma(y, z) \leq 3H + 8\delta$, so

$$(y|z)_s \geq \frac{1}{2}(R - \delta + R - (4H + 7\delta) - (4H + 8\delta)) = R - (4H + 8\delta).$$

Lemma 2.1 allows us to conclude $(\gamma_s^{+\infty} | \gamma^{+\infty})_s \geq R - (4H + 10\delta)$, as desired. □

3. An equivariant map from $X \times \partial G$ to itself

The first step in the proof of Theorem 1.1 is the following construction, which can be thought of as a generalization of that in [BM19, Lemma 3.1]. If X is a space with a proper, free, and cocompact action of G , and $\rho: G \rightarrow \text{Homeo}(Y)$ an action of G on a topological space, one can capture the information of this action as the holonomy of a foliated Y -bundle over X/G ; this is simply the quotient of $X \times Y$ by the diagonal action of G . Here, the case of interest to us is when $Y = \partial G = S^n$. The following proposition gives a construction of a “nice” map between the foliated bundles associated to the boundary action ρ_0 and a small perturbation ρ .

The same proof works with any manifold Y in place of S^n , and any two nearby actions of an arbitrary group G on the space, so we state it in this general context, as follows.

Let Y be a metric space, such that $\text{Homeo}(Y)$ is metrizable and locally contractible. For instance, one may take Y to be any compact manifold, in which case local contractibility of $\text{Homeo}(Y)$ follows from Edwards–Kirby [EK71]. Metrizability of $\text{Homeo}(Y)$ has the following easy consequence.

Observation 3.1. *Let W be a neighborhood of the identity in $\text{Homeo}(Y)$, and let $F \subset \text{Homeo}(Y)$ be finite. Then there is a neighborhood $V \subset W$ of the identity so that the union*

$$\bigcup_{f \in F} fVf^{-1}V$$

lies in W .

Using this, we prove the following.

Proposition 3.2. *Let G be a group, Y a metric space as above, and fix an action $\rho_0: G \rightarrow \text{Homeo}(Y)$. Let X be a metric space on which G acts properly, freely, and cocompactly by isometries.*

For any compact $K \subset X$ and $\epsilon > 0$, there is a neighborhood U of ρ_0 , so that for each $\rho \in U$, there is a homeomorphism $f^\rho: X \times Y \rightarrow X \times Y$ with the following properties:

1. (Covers id_X) *If π_X is the projection from $X \times Y$ to X , then $\pi_X \circ f^\rho = \pi_X$. In other words, f^ρ covers the identity on X .*
2. (Equivariance) *For every $g \in G$, we have*

$$f^\rho(g \cdot x, \rho(g) \cdot \theta) = (g, \rho_0(g)) \cdot f^\rho(x, \theta).$$

3. (Near flatness.) *For any $\theta \in Y$, we have*

$$f^\rho(K \times \{\theta\}) \subset X \times B_\epsilon(\theta).$$

Proof of Proposition 3.2. Since the action is proper, free, and cocompact, there is some $r > 0$ so that every nontrivial element of G moves every point of X a distance at least r . Choose a G -equivariant locally finite cover $\mathcal{U} = \{U_i \mid i \in I\}$ of X by open balls of radius $r/3$, and let N be the nerve of \mathcal{U} . Since \mathcal{U} was G -equivariant and locally finite, the group G acts cocompactly on the simplicial complex N . Since any $g \in G - \{1\}$ moves every set U_i off of itself, G acts freely on N .

We choose a G -equivariant partition of unity $\{\phi_i: X \rightarrow [0, 1] \mid i \in I\}$ subordinate to the cover \mathcal{U} . This partition determines a proper G -equivariant map $\psi: X \rightarrow N$.

To define f^ρ , we first define a map $\varphi^\rho: N \rightarrow \text{Homeo}(Y)$, and then define

$$f^\rho(x, \theta) = (x, (\varphi^\rho \circ \psi(x))(\theta)). \tag{9}$$

The map φ^ρ will be G -equivariant with respect to the “mixed” left action of G by homeomorphisms of $\text{Homeo}(Y)$ given by

$$g \cdot h = \rho_0(g)h\rho(g^{-1}). \tag{10}$$

Definition of φ^ρ .

The definition of φ^ρ is designed to keep track of the compact set K and constant $\epsilon > 0$ for the near flatness condition in the proposition. Let D be a connected union of open simplices in N which meets every G -orbit exactly once. Let K be a compact subcomplex of N , which we assume contains the closed star of any cell of \bar{D} (note that any compact $C \subset X$ has $\psi(C) \subset K$ for some such complex). Let $\epsilon > 0$. Let S be the (finite) set of group elements s so that sD meets the closed star of some vertex in \bar{D} . Let F be the (still finite) set of group elements g , so that $gD \cap K$ is nonempty.

Letting $W = N_\epsilon(\text{id})$, we choose a neighborhood V as in Observation 3.1 so that $\rho_0(g)V\rho_0(g)^{-1}V$ lies in W for all $g \in F$. Now let $m = \dim(N)$. Apply Observation 3.1 and local contractibility of $\text{Homeo}(Y)$ to choose a nested sequence of contractible neighborhoods of 1 inside V :

$$V_0 \subset V_1 \subset \dots \subset V_m \subset V$$

so that for each i and all $s \in S$, we have

$$\rho_0(s)V_i\rho_0(s)^{-1}V_i \subset V_{i+1}. \tag{11}$$

By taking ρ sufficiently close to ρ_0 , we may assume that $\rho_0(g)\rho(g^{-1})$ lies in V_0 for all $g \in F$. We define φ^ρ inductively over the k -skeleta of N in such a way that $\varphi^\rho(\sigma) \subset V_k$ for every k -cell in the closed star of a vertex of \overline{D} .

0-skeleton. Define φ^ρ on $N^{(0)}$ as follows. If $v = gv_0$ for some $v_0 \in D$, then

$$\varphi^\rho(v) = \rho_0(g)\rho(g^{-1}).$$

If v lies in the closed star of some cell of D , then $g \in S$, so $\varphi^\rho(v) \in V_0$ as desired.

Inductive step. Suppose that φ^ρ has been defined on all $(k - 1)$ -cells, and let σ be a k -cell. We may write $\sigma = g\sigma_0$, where σ_0 is an open k -cell in D . The map φ^ρ has already been defined on the boundary of σ_0 , and sends this boundary into V_{k-1} by induction. Using the contractibility of V_{k-1} , we extend φ^ρ over σ_0 in such a way that $\varphi^\rho(\sigma_0) \subset V_{k-1}$. We define $\varphi^\rho|_{\sigma(x)} = \rho_0(g)\varphi^\rho(g^{-1}(x))\rho(g^{-1})$. Since the action on N is free, there is no ambiguity in this definition.

Now suppose that σ lies in the closed star of some vertex of \overline{D} , so $\sigma = s\sigma_0$ for some $s \in S$. The set $\varphi^\rho(\sigma)$ lies in

$$\rho_0(s)V_{k-1}\rho(s^{-1}) = \rho_0(s)V_{k-1}\rho_0(s)^{-1}\rho_0(s)\rho(s)^{-1} \subset \rho_0(s)V_{k-1}\rho_0(s)^{-1}V_0,$$

which lies in V_k by (11). Having verified the inductive hypothesis, we see that we can continue until we have defined φ^ρ equivariantly on all of N . Moreover, we have defined it so that $\varphi^\rho(\sigma_0)$ lies in the neighborhood V for any σ_0 meeting D .

Properties of f^ρ

Having defined φ^ρ , we define f^ρ as in (9).

$$f^\rho(x, \theta) = (x, (\varphi^\rho \circ \psi)(\theta)).$$

By definition, this covers the identity map on X . To simplify notation, let Φ^ρ denote $\varphi^\rho \circ \psi$. Note that Φ^ρ satisfies equivariance as φ^ρ does. This also gives equivariance of f^ρ , as follows:

$$\begin{aligned} f^\rho(g \cdot x, \rho(g) \cdot \theta) &= (g \cdot x, (\Phi^\rho(g \cdot x)\rho(g)) \cdot \theta) \\ &= \left(g \cdot x, \left(\rho_0(g)\Phi^\rho(x)\rho(g^{-1})\rho(g) \right) \cdot \theta \right) \\ &= (g \cdot x, (\rho_0(g)\Phi^\rho(x)) \cdot \theta) \\ &= (g, \rho_0(g)) \cdot (x, \Phi^\rho(x) \cdot \theta) \\ &= (g, \rho_0(g)) \cdot f^\rho(x, \theta). \end{aligned}$$

It remains to check near flatness. For any cell σ of the larger compact complex K , there is some $g \in F$ and some $\sigma_0 \subset D$ so that $\sigma = g\sigma_0$. Equivariance tells us that

$$\begin{aligned} \varphi^\rho(\sigma) &= \rho_0(g)\varphi^\rho(\sigma_0)\rho(g^{-1}) \\ &= \rho_0(g)\varphi^\rho(\sigma_0)\rho_0(g)^{-1} \cdot \rho_0(g)\rho(g^{-1}) \\ &\subset \rho_0(g)V\rho_0(g)^{-1} \cdot V \subset W. \end{aligned}$$

Since every homeomorphism in W moves every point of Y a distance of at most ϵ , we have $f^\rho(\sigma \times \{\theta\}) \subset B_\epsilon(\theta)$ as desired. □

4. Reduction to the main case

In this section, we reduce to the case $\partial G = S^n$ for $n \geq 2$, and also to the case where G acts faithfully on its boundary. The first reduction (to $n \geq 2$) comes from combining work of Matsumoto with the Convergence Group Theorem.

Proposition 4.1. *Let G be a hyperbolic group with circle boundary. Then the action of G on ∂G is topologically stable.*

This can quite likely be derived from Matsumoto’s original proof, as the main techniques are Euler characteristic and lifting to covers. For completeness, we give a short argument using standard tools from circle dynamics.

Proof. Let G be a hyperbolic group with circle boundary. By the Convergence Group Theorem [Gab92, CJ94], there exists a normal, finite index torsion-free subgroup G' of G that is isomorphic to the fundamental group of a closed surface. Let ρ be a perturbation of the standard boundary action ρ_0 of G . By [Mat87], there exists a continuous, surjective, degree one map $h : S^1 \rightarrow S^1$, such that $h\rho(g) = \rho_0(g)h$ for each $g \in G'$. If the action of G' is minimal, then h is a conjugacy. If the action of G' is not minimal, then there exists a unique invariant *exceptional minimal set* X , homeomorphic to a Cantor set, and h collapses the closure of each complementary interval to a point and is otherwise injective (see e.g. [Ghy01, Proposition 5.6]). It is also easy to see that h varies continuously with ρ , so can be taken as close to the identity as desired by taking ρ close to ρ_0 .

Since G' is normal in G , the set X is $\rho(G)$ -invariant. It follows that G permutes the point-preimages of h . From this, we will now deduce that h in fact defines a semiconjugacy intertwining the actions of $\rho_0(G)$ and $\rho(G)$. To see this, we use the fact that attracting fixed points of elements of $\rho_0(G')$ are dense in S^1 . If x is the attracting fixed point of $\rho_0(\gamma)$ for some $\gamma \in G'$, then $\rho_0(g)x$ is the attracting fixed point of $\rho_0(g\gamma g^{-1})$, an element which also lies in G' . Thus, for any $y \in h^{-1}(x)$, we have $\rho(g)(y) \in h^{-1}\rho_0(g)(x)$; equivalently, $h\rho(g)(y) = \rho_0(g)h(y)$. Since h is continuous, and the union of preimages of attracting fixed points is dense in S^1 , this shows that $h\rho(g) = \rho_0(g)h$ holds globally. □

Proposition 4.2. *Suppose G is a hyperbolic group with sphere boundary, and let $F < G$ be the subgroup of elements which act trivially on ∂G . Then G is topologically stable if and only if G/F is topologically stable.*

Proof. Since F is a finite normal subgroup, the canonical action of G on its boundary factors through the canonical action of G/F on its boundary, and these boundaries are the same. Let N be the maximum order of an element of F . By a theorem of Newman [New31], there is a neighborhood U of the identity in $\text{Homeo}(S^n)$ so that any torsion element in U has order greater than N . Thus, for any sufficiently small perturbation of the canonical action, the elements of F will still act trivially, and so small perturbations of the canonical action of G on its boundary are in one-to-one correspondence with small perturbations of G/F on its boundary. □

We therefore make the following assumptions for the remainder of the paper.

Assumption 4.3. *The hyperbolic group G acts faithfully on its boundary ∂G .*

Assumption 4.4. *The boundary of G is a topological sphere of dimension at least two.*

5. A space with a proper, cocompact, and free action of G

Recall that Ξ is the space of distinct triples in ∂G . Since ∂G is assumed to be a sphere of dimension $n \geq 2$ (Assumption 4.4), the triple space Ξ is a connected $3n$ -manifold. The following is an easy consequence of [AMN11, Theorem 1.1].

Proposition 5.1. *There is a proper G -invariant metric on Ξ .*

We fix such a metric d_Ξ now. If additionally no nontrivial element of G fixes more than two points of ∂G , then the action of G on Ξ is free and Ξ can play the role of the metric space X in Proposition 3.2. In general, G may not act freely on Ξ , so a different space is needed. In this section, we show how to build such a space with a free action by deleting a small, G -equivariant neighborhood of the set of points in Ξ with nontrivial stabilizer.

Notation 5.2. For $g \in G$, we denote by $\text{fix}(g) \subset \partial G$ the set of points fixed by the natural action of g . We set

$$F = \{(a, b, c) \in \Xi \mid \{a, b, c\} \subset \text{fix}(g) \text{ for some nontrivial } g \in G\}.$$

Recall a subset S of a metric space M is ϵ -dense if every point of M is distance at most ϵ of a point of S . Also recall that $N_\kappa(S)$ denotes the open κ -neighborhood of a set S . If S is a subset of Ξ , this neighborhood is to be taken with respect to d_Ξ . Our first goal is to establish the following.

Proposition 5.3 (F is sparse). *For any $\epsilon > 0$, there exists $\kappa > 0$, such that $\Xi - N_\kappa(F)$ is ϵ -dense in Ξ .*

Since we will frequently refer to $\Xi - N_\kappa(F)$, we fix the following notation.

Notation 5.4. For $\kappa > 0$, we let $X_\kappa = \Xi - N_\kappa(F)$.

Proposition 5.3 has the following useful corollary.

Corollary 5.5. *For sufficiently small κ , property ($\delta 4$) of Lemma 2.7 still holds, with π_δ replaced by the restriction of π_δ to X_κ .*

Proof. By equivariance and the fact that $\pi(x)$ is nonempty for every triple in Ξ , Property ($\delta 4$) holds as soon as X_κ is nonempty. □

The proof of Proposition 5.3 requires several preliminary results, starting with the following.

Lemma 5.6. *For any torsion element $g \in G$, such that $\text{fix}(g)$ has at least three points, there exists a quasi-convex subgroup $Q_g \subset G$, such that $\Lambda(Q_g) = \text{fix}(g)$.*

Here we use the standard notation $\Lambda(Q_g)$ for the limit set of Q_g in ∂G .

Proof. This proof is adapted from an argument of Kapovich [Kap]. Fix a torsion element g . Let H be a maximal torsion subgroup pointwise fixing $\text{fix}(g)$, and let Q be the normalizer of H in G . We claim Q is the stabilizer of $\text{fix}(H) = \text{fix}(g)$. That Q preserves this set is immediate. For the reverse inclusion, if $f \in G$ preserves $\text{fix}(H)$, then fHf^{-1} pointwise fixes $\text{fix}(H)$ as well. Since $\text{fix}(H)$ has at least three points in it, the subgroup generated by H and fHf^{-1} cannot contain a loxodromic, and is therefore finite [GdlH90, Chapter 8, Section 3]. By maximality of H , we have that $f \in Q$. This shows Q is the stabilizer of $\text{fix}(g)$.

To conclude the proof, we wish to show that Q is quasi-convex and $\Lambda(Q) = \text{fix}(g)$. We can then take $Q_g = Q$ in the conclusion. Let C denote the quasi-convex hull of $\text{fix}(g)$, meaning the set of all biinfinite geodesics with both endpoints in G . Note that there is a uniform bound, say r' , on distance that any $h \in H$ can translate any point in C .

We now show Q acts cocompactly on C , which is enough to show Q has the desired properties. To see this, assume for contradiction that C contains infinitely many distinct cosets $\{Qg_k\}$. Since every $h \in H$ translates each g_n a distance at most r' , the conjugates $g_k^{-1}Hg_k$ all lie in the r' -ball about the

identity in Γ . It follows that for infinitely many pairs i, j , the subgroups $g_i^{-1}Hg_i = g_j^{-1}Hg_j$ agree, and thus $g_i g_j^{-1} \in Q$. We conclude the cosets $\{Qg_k\}$ were not distinct, giving the desired contradiction. \square

Recall that a *null sequence* in a metric space is a collection of subsets \mathcal{D} so that for every $\epsilon > 0$, the set $\{D \in \mathcal{D} \mid \text{diam}(D) > \epsilon\}$ is finite.

Corollary 5.7. *The collection of fixed point sets of nontrivial torsion elements is a null sequence in ∂G .*

Proof. Since G is hyperbolic, it contains only finitely many conjugacy classes of torsion elements (see [Gro87, 2.2 B]). It therefore suffices to show the fixed point sets of elements in a single conjugacy class form a null sequence. These fixed point sets are exactly the G -translates of $\text{fix}(g)$ for some torsion element g . By Lemma 5.6, $\text{fix}(g)$ is equal to $\Lambda(Q_g)$ for a quasi-convex subgroup Q_g of G . By Assumption 4.3, g does not fix all of ∂G , so this subgroup Q_g must be infinite index in G . By [GMRS98, Corollary 2.5], the G -translates of $\Lambda(Q_g)$ form a null sequence. \square

Proposition 5.8. *F is a closed set with empty interior.*

Proof. Let $K \subset \Xi$ be an arbitrary compact set. By Corollary 5.7, there are only finitely many elements g , so that $\text{diam}(\text{fix}(g)) \geq \text{minsep}(K)$. There are therefore only finitely many nontrivial g , so that K contains a triple of points in $\text{fix}(g)$. For each such g , the set $\text{fix}(g) \times \text{fix}(g) \times \text{fix}(g)$ intersects K in a closed subset with empty interior in Ξ . It follows that $F \cap K$ is closed with empty interior. Since Ξ can be exhausted by compact sets, the conclusion follows. \square

We also need the following general result about closed sets with empty interior in compact metric spaces.

Lemma 5.9. *Let A be a compact metric space and $C \subset A$ a closed subset with empty interior. Given any $\epsilon > 0$, there exists $\kappa > 0$, so that $A - N_\kappa(C)$ is ϵ -dense in A .*

Proof. Let d_A denote the metric on A . Since C has empty interior, for all $x \in A$, there exists a point $p_x \notin C$ with $d_A(p_x, x) < \epsilon$. Since A is compact, there is a finite collection x_1, \dots, x_k so that the open ϵ -balls $B_\epsilon(p_{x_i})$ cover A . The set C is closed, so the distance $d_A(p_{x_i}, C)$ is positive for each i . We let κ be half the minimum of the distances $d_A(p_{x_i}, C)$. Each $x \in A$ is contained in one of the balls $B_\epsilon(p_{x_i})$, so the set $\{p_{x_1}, \dots, p_{x_k}\} \subset A - N_\kappa(C)$ is ϵ -dense. \square

Proof of Proposition 5.3. By Proposition 5.8, F is closed with empty interior. Recall from Proposition 2.2 that G acts properly discontinuously and cocompactly on Ξ . Thus, $F/G \subset \Xi/G$ is closed with empty interior, and the proposition follows immediately from Lemma 5.9, taking $A = \Xi/G$ and $C = F/G$. \square

5.1. Additional properties of X_κ

We establish some properties of the sets $X_\kappa := \Xi - N_\kappa(F)$. First, we make the following observation, relevant to the application of Proposition 3.2.

Lemma 5.10. *For any $\kappa > 0$, the group G acts properly, freely, and cocompactly by isometries on X_κ .*

Proof. The group G already acts properly and cocompactly on X , and $N_\kappa(F)$ is open and G -invariant, so G still acts properly and cocompactly on X_κ . The only points of Ξ with nontrivial stabilizer are in F , which has been removed. \square

Next, we give a technical refinement of Proposition 5.3, which is what we actually use in the next section.

Lemma 5.11. *Suppose we are given a compact set $D \subset \Xi$, and $\epsilon'' < \epsilon' < \text{minsep}(D)/2$. There exists $\kappa > 0$, such that, for any $(a, \theta, c) \in D$, there exists c' so that $d_{\text{vis}}(c, c') < \epsilon''$ and $\{a\} \times B_{\epsilon'}(\theta) \times \{c'\}$ lies in the interior of X_κ .*

We recall that minsep is defined in terms of the visual metric d_{vis} .

Proof. Let D' be a compact subset of Ξ containing $B_{\epsilon'}(x) \times B_{\epsilon'}(y) \times B_{\epsilon'}(z)$ for all $(x, y, z) \in D$. Let F be the set of elements $g \in G - \{e\}$ so that g fixes some triple in D' . By Corollary 5.7, the fixed point sets of torsion elements form a null sequence in ∂G . For each $g \in F$, we have $\text{diam}(\text{fix}(g)) \geq \text{minsep}(D') > 0$, so F is finite. Let $C = \bigcup \{\text{fix}(g) \mid g \in F\}$ be the union of these fixed sets in ∂G . Since C is a finite union of closed sets with empty interior, C is closed with empty interior.

By Lemma 5.9, we can choose a positive $\lambda < \epsilon''$, such that $\partial G - N_\lambda(C)$ is ϵ'' -dense in ∂G . Choose κ small enough so that $N_{2\kappa}(F) \cap D'$ is a subset of the neighborhood of F of radius λ in the product visual metric on $(\partial G)^3$.

To verify this κ works, choose $(a, \theta, c) \in D$. Since $\partial G - N_\lambda(C)$ is ϵ'' -dense, we can find $c' \in \partial G - N_\lambda(C)$ with $d_{\text{vis}}(c, c') < \epsilon''$. We claim that $\{a\} \times B_{\epsilon'}(\theta) \times \{c'\}$ is contained in $\Xi - N_{2\kappa}(F)$, and hence in the interior of X_κ . Indeed, for any $b \in B_{\epsilon'}(\theta)$, we have $(a, b, c') \in D'$. Suppose $(x, y, z) \in F$ is a closest point of F to (a, b, c') in the product metric. If the visual distance between each coordinate of (x, y, z) and (a, b, c') were less than λ , which is less than ϵ' , then we would have $(x, y, z) \in D' \cap F$. By our definition of C , this means that x, y , and z all belong to C . But the minimum distance from c' to a point of C is greater than λ by construction, so we conclude (a, b, c') lies outside the λ -neighborhood of F in the product metric, and hence outside $N_{2\kappa}(F)$. □

Finally, we need a technical result which will come into play at the very end of the proof, when we show our semiconjugacy is well-defined. To state it, we need a definition.

Definition 5.12 (κ path property). We say a compact $D \subset \Xi$ has the κ path property if the following holds:

For each $b \in \partial G$, and pair of points (a, b, c) and (a', b, c') in D , there is a path a_t from a to a' and $c'' \in \partial G$, such that $(a_t, b, c'') \in D - N_\kappa(F)$ for all t .

Lemma 5.13. *For any compact set $K \subset \Xi$, there exists $\kappa > 0$ and a compact set $D \supset K$ so that D has the κ path property.*

Proof. Fix a round metric d_{rnd} on $\partial G = S^n$. The collection

$$C_m := \left\{ (a, b, c) \in \Xi \mid d_{\text{rnd}}(a, b) \geq \frac{1}{m}, d_{\text{rnd}}(a, c) \geq \frac{1}{m}, d_{\text{rnd}}(b, c) \geq \frac{1}{m} \right\}$$

forms an exhaustion of Ξ by compact sets, so we may choose m sufficiently large so that $K \subset C_m$ and the diameter of S^n in the round metric is larger than $\frac{10}{m}$.

Since D is compact, only finitely many elements of G fix a triple of points that meets D . Let $C \subset S^n$ denote the union of these finitely many fixed sets; it is a closed subset of ∂G with empty interior.

By Lemma 5.9, there exists $\lambda < \frac{1}{m}$, such that $S^n - N_\lambda(C)$ is $\frac{1}{m}$ -dense in S^n (again, using the round metric). Consider a shortest length geodesic path with respect to the round metric on S^n between a and a' . If this path does not meet the closed $\frac{1}{m}$ -ball about b , then call this path a_t . Otherwise, modify this path by removing the segment that intersects the closed $\frac{1}{m}$ -ball, replacing it with a path that lies on the boundary of this ball, and call the resulting path a_t (we are here using our Assumption 4.4 that $n \geq 2$). In either case, a_t lies in the union of a radius $\frac{1}{m}$ -ball and a geodesic segment.

Since the diameter of S^n in the chosen round metric is larger than $\frac{10}{m}$, we may find a point $\hat{c} \in S^n$ that avoids both the $\frac{2}{m}$ -neighborhood of this path a_t and the $\frac{2}{m}$ -neighborhood of b . Since $S^n - N_\lambda(C)$ is $\frac{1}{m}$ -dense, we may move this point a distance of at most $\frac{1}{m}$ to find a point $c'' \notin N_\lambda(C)$. Thus, we have $(a_t, b, c'') \in D$ for all t .

Finally, choose κ small enough so that any element of a triple in $N_\kappa(F) \cap D$ lies within the λ -neighborhood (in the round metric) of some point of C . Note that this choice of κ only depends

on D and λ , which both depended only on the set K . The point c'' was chosen so that $c'' \notin N_\lambda(C)$. Thus, we conclude that $(a_t, b, c'') \notin N_\kappa(F)$ holds for all t , as desired. \square

Note that even though d_{vis} and the round metric give the same topology on S^n , the metric d_{vis} may be rather strange-looking. In particular, d_{vis} is not Riemannian, so even very small balls with respect to d_{vis} may not be connected. This is the reason for the use of the round metric in the proof above.

6. Nearby representations give quasi-geodesic partitions

In this section, we fix a constant κ to define a space $X_\kappa \subset \Xi$ as in Section 5, and fix a neighborhood of ρ_0 . Our goal is to show that every representation ρ in this neighborhood has ρ_0 as a factor. In Section 6.2, we show that representations in this neighborhood induce quasi-geodesic partitions of a section of $X_\kappa \times \partial G \rightarrow X_\kappa$. The endgame of the proof, carried out in Section 7, consists of using the endpoints of these quasi-geodesics to build a semiconjugacy.

6.1. Fixing a neighborhood of ρ_0

We first fix a neighborhood V of the identity in the set of continuous self-maps $\partial G \rightarrow \partial G$. By shrinking V if necessary, we may assume that V consists of degree one maps. There is some $C_V > 0$ so that any map h satisfying the lower bound on Gromov products

$$\forall x \in \partial G, \quad (h(x) | x)_e > C_V \tag{\dagger}$$

lies in V .

In Definition 6.4, we will specify a neighborhood $\mathcal{U}(V)$ of ρ_0 in $\text{Hom}(G, \text{Homeo}(\partial G))$, ultimately showing that any $\rho \in \mathcal{U}(V)$ is semiconjugate to ρ_0 by a map h satisfying (\dagger) . This is sufficient to prove Theorem 1.1.

To specify $\mathcal{U}(V)$ and set up the proof, we need to fix several intermediate constants and compact sets. Recall \mathcal{S} denotes our chosen generating set for G .

Notation 6.1. Let $H = \max\{2\delta, Q(3\delta)\} + 1$, where $Q(\cdot)$ is from Lemma 2.5.

Lemma 6.2. *There exist compact sets $D_0 \subset D_{\frac{1}{2}} \subset D_1 \subset D_2$ in Ξ and positive constants $\epsilon_2 \leq \epsilon_1$ and κ_0 satisfying the following.*

1. D_0 contains $\pi^{-1}(\mathbf{e})$ (hence, $GD_0 = \Xi$).
2. $D_{\frac{1}{2}} = \bigcup \{gD_0 \mid g \in \mathcal{S} \cup \mathcal{S}^{-1} \cup \{\mathbf{e}\}\}$.
3. $\epsilon_1 < \frac{1}{2} \text{minsep}(D_{\frac{1}{2}})$.
4. D_1 contains $B_{\epsilon_1}(x) \times B_{\epsilon_1}(y) \times B_{\epsilon_1}(z)$ for every $(x, y, z) \in D_{\frac{1}{2}}$.
5. D_1 contains $\pi^{-1}(B_R(\mathbf{e}))$, where

$$R > \max\{24H + 52\delta + \text{diam}(\pi(D_0)), C_V + 4H + 11\delta\}. \tag{\star\star}$$

6. For all $\kappa < \kappa_0$, the κ path property of Lemma 5.13 holds for D_1 ,
7. $\epsilon_2 < \frac{1}{2} \text{minsep}(D_1)$.
8. D_2 contains $B_{\epsilon_2}(x) \times B_{\epsilon_2}(y) \times B_{\epsilon_2}(z)$ for every $(x, y, z) \in D_1$.

The sets D_0, D_1 , and D_2 play a major role in the rest of the proof. The set $D_{\frac{1}{2}}$ will only appear in the last step.

Proof. We let D_0 be a compact subset of Ξ containing $\pi^{-1}(B_1(\mathbf{e}))$ and define $D_{\frac{1}{2}}$ as in Item (2). We fix any positive ϵ_1 satisfying Item (3). Let $K \subset \Xi$ be a compact set large enough to contain the union of $\pi^{-1}(B_R(\mathbf{e}))$, where R is in equation $(\star\star)$, as well as the union of the sets $B_{\epsilon_1}(x) \times B_{\epsilon_1}(y) \times B_{\epsilon_1}(z)$ for every $(x, y, z) \in D_{\frac{1}{2}}$. Applying Lemma 5.13 to K gives us a constant κ_0 and set D_1 containing K so

that the κ_0 path property holds for D_1 ; note that by definition, the κ path property remains true for any $\kappa < \kappa_0$, since $N_\kappa(F) \subset N_{\kappa_0}(F)$. Then take any ϵ_2 and D_2 as described. \square

Lemma 6.3. *There exists positive $\epsilon < \frac{1}{2}\epsilon_2$, such that the following hold:*

- ($\epsilon 1$) *Any geodesic between points a, b with $d_{\text{vis}}(a, b) \leq 2\epsilon$ has distance at least 10δ from $\pi(D_2)$ (and hence from $B_R(\mathbf{e})$).*
- ($\epsilon 2$) *For every $p \in \partial G$, there is a closed contractible set B_p so that*

$$B_\epsilon(p) \subset B_p \subset B_{\epsilon_2}(p).$$

Proof. That the first condition can be met is a consequence of compactness. Indeed, by Lemma 2.6, the 10δ -neighborhood of the set $\pi(D_2)$ is bounded in Γ , hence contained in some compact ball centered at \mathbf{e} . Thus, there is a positive lower bound on the visual distance between the endpoints of any geodesic passing through that ball.

We now address the second condition. Since S^n is compact, the visual and round metrics are uniformly equivalent. Thus, there is a $\mu > 0$ so that the round ball of radius 2μ about p is contained in $B_{\epsilon_2}(p)$ for every $p \in S^n$, so we can take B_p to be the closed round μ -ball. There is then some $\epsilon > 0$ so that $B_\epsilon(p) \subset B_p$ for every p . \square

Definition 6.4 (The neighborhood $\mathcal{U}(V)$). Recall we fixed $C_V > 0$ at the beginning of the subsection, see Equation (\dagger). Using this C_V , fix the compact sets $D_0 \subset D_{\frac{1}{2}} \subset D_1 \subset D_2 \subset \Xi$ and positive constants $\epsilon < \epsilon_2 \leq \epsilon_1$ and κ_0 satisfying the conclusions of Lemmas 6.2 and 6.3 above.

Fix some $\kappa > 0$ satisfying

- 1. $\kappa < \kappa_0$; and
- 2. For any $(a, c, \theta) \in D_1$, there exists $c' \in B_\epsilon(c)$ so that $\{a\} \times B_{\epsilon_2}(\theta) \times \{c'\}$ lies in the interior of X_κ .

That such a κ exists follows by applying Lemma 5.11 with $D = D_1$, $\epsilon' = \epsilon_2$, and $\epsilon'' = \epsilon$. Note that $\epsilon < \epsilon_2 < \text{minsep}(D_1)/2$, so the hypotheses of that lemma are satisfied.

Define $\mathcal{U}(V)$ to be a neighborhood of ρ_0 in $\text{Hom}(G, \text{Homeo}(\partial G))$ consisting of representations satisfying both of the following:

- 3. $d_{\text{vis}}(\rho(g), \rho_0(g)) < \epsilon$ for every $g \in \mathcal{S} \cup \mathcal{S}^{-1}$; and
- 4. the map f^ρ defined in Section 3, taking $X = X_\kappa$, has the property that $f^\rho((D_2 \cap X) \times \{\theta\}) \subset B_\epsilon(\theta)$ for every θ .

We now fix notation.

Convention 6.5. *Fix some $\rho \in \mathcal{U}(V)$. Since ρ is fixed, we henceforth drop it from the notation, writing $f = f^\rho$. Since κ is also fixed, we also drop it from our notation when convenient, writing X for X_κ .*

We keep this convention for the remainder of the work. Our eventual goal is to show that ρ_0 is a factor of ρ via a semiconjugacy satisfying (\dagger). The reader will note that when $F = \emptyset$, then $X = \Xi$ (for example, this holds when G is torsion-free). It may be helpful on a first reading to keep this special case in mind, thinking of X as the space of distinct triples.

6.2. A section partitioned into coarsely geodesic sets

We begin by describing a natural foliation on a section of $\Xi \times \partial G \rightarrow \Xi$. Let $\sigma: \Xi \rightarrow (\Xi \times \partial G)$ be the section given by $\sigma((a, b, c)) = ((a, b, c), b)$. The image of σ has a topological foliation by leaves

$$L_a := \{((a, b, c), b) \mid b, c \in \partial G - \{a\}, b \neq c\}.$$

The leaves of this foliation are *coarsely geodesic* in the following sense. For any $a \neq \theta$, the set $L_a \cap (\Xi \times \{\theta\})$ is the image in $\sigma(\Xi)$ of a set whose projection to Γ is Hausdorff distance at most δ from any geodesic joining a and θ . The set $L_a \cap (\Xi \times \{a\})$ is empty.

We will show that the sets

$$L_a \cap f^\rho(X \times \{\theta\}) \subset \sigma(X)$$

behave much like the sets $L_a \cap (\Xi \times \{\theta\})$, in the sense that for each a and θ in ∂G , this set is either empty or looks “coarsely geodesic,” meaning the projection of each nonempty set $L_a \cap f^\rho(X \times \{\theta\})$ to the Cayley graph Γ lies at a uniformly bounded Hausdorff distance from a biinfinite geodesic in Γ .

To formalize this, define $\bar{\pi}$ on $\sigma(X)$ by $\bar{\pi} = \pi \circ \sigma^{-1}$. It follows from the definition of π on sets that for $Z \subset \sigma(X)$, the set $\bar{\pi}(Z)$ is equal to the union $\bigcup_{z \in Z} \bar{\pi}(z)$. Our goal for the section can now be restated:

Proposition 6.6. *For any $(a, \theta) \in \partial G^2$, either $L_a \cap f^\rho(X \times \{\theta\})$ is empty, or $\bar{\pi}(L_a \cap f^\rho(X \times \{\theta\}))$ is uniformly bounded Hausdorff distance from a biinfinite geodesic in Γ with one endpoint equal to a .*

In fact, we eventually prove something slightly stronger (Proposition 6.11), but the statement and proof require some set-up.

Notation 6.7. For $\theta \in \partial G$, we let $Y_\theta = f(X \times \{\theta\}) \subset X \times \partial G$. For $a, c, \theta \in \partial G$, we let $f_{a,c,\theta}$ denote the map $x \mapsto f_\theta(a, x, c)$. This map is defined on all points $x \in \partial G$, such that $(a, x, c) \in X$ and it is continuous on its domain of definition.

Recall $L_a \cap (\Xi \times \{\theta\})$ is nonempty provided $a \neq \theta$. Our first lemma ensures that $L_a \cap Y_\theta$ is nonempty provided a and θ are far apart, as follows:

Lemma 6.8. *For any $(a, \theta, c) \in D_1$, the set $L_a \cap Y_\theta$ is nonempty, and contains a point of the form $((a, b_\theta, c'), b_\theta)$, where $b_\theta \in B_\epsilon(\theta)$, and $c' \in B_\epsilon(c)$; in particular, we have $(a, b_\theta, c') \in D_2 \cap X$.*

Proof. By Item (2) of Definition 6.4, there exists a point $c' \in B_\epsilon(c)$, such that $\{a\} \times B_{\epsilon_2}(\theta) \times \{c'\} \subset X$. Thus, $f_{a,c',\theta}$ is defined on $B_{\epsilon_2}(\theta)$ and Definition 6.4, Item (4) implies that $f_{a,c',\theta}(B_{\epsilon_2}(\theta)) \subset B_\epsilon(\theta)$. Lemma 6.3 states that there exists a closed, contractible set B_θ , such that

$$B_\epsilon(\theta) \subset B_\theta \subset B_{\epsilon_2}(\theta).$$

Thus,

$$f_{a,c',\theta}(B_\theta) \subset f_{a,c',\theta}(B_{\epsilon_2}(\theta)) \subset B_\epsilon(\theta) \subset B_\theta.$$

By the Lefschetz fixed point theorem, $f_{a,c',\theta}$ has a fixed point $b_\theta \in B_\epsilon(\theta)$, which means exactly that $((a, b_\theta, c'), b_\theta) \in L_a \cap Y_\theta$. Recall that $\epsilon < \epsilon_2$, so by our choice of sets and constants as in Lemma 6.2, we have also $(a, b_\theta, c') \in D_2$. □

Notation 6.9. Going forward, we define $S(a, \theta) := \bar{\pi}(L_a \cap Y_\theta) \subset \Gamma$.

Our goal is to show that, whenever $S(a, \theta)$ is nonempty, it is bounded Hausdorff distance from a biinfinite geodesic with one endpoint equal to a . The following lemma gives a local estimate. Recall we have fixed $H = \max\{2\delta, Q(3\delta)\} + 1$, where $Q(\cdot)$ is the function from Lemma 2.5.

Lemma 6.10. *Suppose there is a point $((a, b_0, c_0), b_0)$ in $\sigma(D_1) \cap L_a \cap Y_\theta$. Let γ be a geodesic from a to b_0 in Γ . Then*

$$S(a, \theta) \cap B_R(\mathbf{e}) \subset N_H(\gamma), \text{ and } \gamma \cap B_R(\mathbf{e}) \subset N_H(S(a, \theta)).$$

Proof. For the first inclusion, suppose that p lies in $S(a, \theta) \cap B_R(\mathbf{e})$. Then, since D_1 contains $\pi^{-1}(B_R(\mathbf{e}))$, we have $p \in \bar{\pi}(f((D_1 \cap X) \times \{\theta\}) \cap L_a)$. In particular, there exist $b, c \in \partial G$, such that all the following hold simultaneously:

1. $p \in \pi(a, b, c)$,
2. $(a, b, c) \in D_1 \cap X$,
3. $f_\theta(a, b, c) = b$.




Figure 1. $p \in \pi(a, b, c)$ is close to any point q of $\pi(a, b', c')$.

Both (a, b, c) and (a, b_0, c_0) lie in $(D_1 \cap X) \subset (D_2 \cap X)$. By item (4) in Definition 6.4, we conclude that $b = f_\theta(a, b, c)$ and $b_0 = f_\theta(a, b_0, c_0)$ both lie in $B_\epsilon(\theta)$. In particular, $d_{\text{vis}}(b, b_0) < 2\epsilon$, so any geodesic from b to b_0 misses $B_R(\mathbf{e})$ by at least 10δ (Lemma 6.3.($\epsilon 1$)). By δ -slimness of ideal triangles, this implies that if $[a, b]$ is any geodesic joining a to b , then $[a, b] \cap B_{R+8\delta}(\mathbf{e})$ lies in the δ -neighborhood of γ (the geodesic joining a to b_0). By definition of π , the point p lies within δ of a point on $[a, b]$ which is at most $R + \delta$ from \mathbf{e} , so lies within 2δ of γ . This shows the first inclusion.

For the second inclusion, suppose that $p \in \gamma \cap B_R(\mathbf{e})$. By assumption ($\delta 5$) on δ , there is some c so that $p \in \pi(a, b_0, c)$. Since D_1 contains $\pi^{-1}(B_R(\mathbf{e}))$, we have $(a, b_0, c) \in D_1$.

By Item (2) of Definition 6.4, there exists $c' \in B_\epsilon(c)$, such that

$$\{a\} \times B_{\epsilon_2}(b_0) \times \{c'\} \subset (X \cap D_2).$$

Let $B = B_\theta$ be the contractible set from Lemma 6.3.($\epsilon 2$). This set contains $B_\epsilon(\theta)$, so in particular, $b_0 \in B$. We have $\{a\} \times B \times \{c'\} \subset X \cap D_2$, so as in the proof of Lemma 6.8, f_θ induces a map $f_{a,c',\theta}: B \rightarrow B_\epsilon(\theta) \subset B$, and $f_{a,c',\theta}$ fixes some $b' \in B_\epsilon(\theta) \subset B_{2\epsilon}(b_0)$. The point $((a, b', c'), b')$ is therefore in $L_a \cap Y_\theta$, so $\pi(a, b', c')$ is a subset of $S(a, \theta)$. Let $q \in \pi(a, b', c')$.

See Figure 1.

We claim q is close to p . To see this, we first show p is within 2δ of any geodesic from a to b' . Fix such a geodesic $[a, b']$, and geodesics $[b_0, b']$ and $[b_0, a]$. Since $p \in B_R(\mathbf{e})$, it lies distance at least 10δ from any point on $[b_0, b]$. Also, we have that p lies within δ of some point on $[b_0, a]$, so by δ -thinness, p is distance at most 2δ from a point on $[a, b']$.

Repeating this argument with c' in place of b' shows that p is within 2δ of any geodesic from a to c' . A slight modification shows that p is within 3δ of any geodesic $[b', c']$, as follows: Considering a quadrilateral with sides $[b_0, c]$, $[c, c']$, $[b', c']$, and $[b_0, b']$, we know that p is within δ of some point on $[b_0, c]$ and this must lie within 2δ of a point on $[b', c']$ since the other two sides are each distance at least 10δ from p . Thus, $p \in \pi_{3\delta}(a, b', c')$, and of course q lies in this set as well.

By Lemma 2.5, $\pi_{3\delta}(a, b, c)$ has diameter at most $Q(3\delta)$. In particular, $d_\Gamma(p, q) \leq Q(3\delta)$. Since $q \in S(a, \theta)$, this shows the second inclusion. \square

We will now combine this work with Lemma 2.12 to prove Proposition 6.6. We will actually prove the following slightly stronger statement.

Proposition 6.11. *If $S(a, \theta) \neq \emptyset$, then $S(a, \theta)$ is Hausdorff distance less than $3H + 6\delta + 1$ from a geodesic γ with one endpoint at a .*

If in addition $L_a \cap Y_\theta \cap \sigma(D_0)$ is nonempty, then γ joins a to a point e^+ so that $(e^+ | \theta)_\mathbf{e} \geq R - (4H + 11\delta)$.




Figure 2. γ and γ_s should both be close to S_0 on the shaded region $B_R(s)$, giving a contradiction if a is not an endpoint of γ .

Proof. We first verify the hypotheses of Lemma 2.12 for the set $S = S(a, \theta) \cap G$ (recall that G is canonically identified with the vertices of Γ). By Corollary 5.5, π is surjective, so for each $s \in S$, we can choose some b_s, c_s so that $\sigma(a, b_s, c_s) \in L_a \cap Y_\theta$ and $s \in \pi(a, b_s, c_s)$. Since $\pi^{-1}(\mathbf{e}) \subset D_0$, the point $\sigma(a, b_s, c_s)$ lies in $\sigma(s \cdot D_0) \cap L_a \cap Y_\theta$. Left translating by s^{-1} , we have

$$(s^{-1}a, s^{-1}b_s, s^{-1}c_s) \in \sigma(D_0) \cap L_{s^{-1}a} \cap Y_{\rho(s^{-1})\theta}.$$

Let $\hat{\gamma}_s$ be any biinfinite geodesic from $s^{-1}a$ to $s^{-1}b_s$. Lemma 6.10 implies

$$s^{-1}S \cap B_R(\mathbf{e}) \subset N_H(\hat{\gamma}_s), \text{ and } \hat{\gamma}_s \cap B_R(\mathbf{e}) \subset N_H(S).$$

Setting γ_s equal to $s \cdot \hat{\gamma}_s$, we find that

$$S \cap B_R(s) \subset N_H(\gamma_s), \text{ and } \gamma_s \cap B_R(s) \subset N_H(S).$$

Now fix any $\frac{R}{4}$ -connected component S_0 of S so that we may apply Lemma 2.12 (we will see later that S is $\frac{R}{4}$ -connected, so in fact $S_0 = S$). The paragraph above shows that all the hypotheses of Lemma 2.12 hold for S_0 , so we conclude that there is a biinfinite geodesic γ with $d_{\text{Haus}}(\gamma, S_0) \leq 3H + 6\delta$.

We next claim that one of the endpoints of γ is a . We argue by contradiction. Using δ -slimness of ideal triangles, there is a point $p \in \gamma$ so that p is within 2δ of any geodesic joining a to any endpoint of γ . Since S_0 is Hausdorff distance at most $3H + 6\delta$ from γ , there is some $s \in S_0$ so that $d_\Gamma(s, p) \leq 3H + 6\delta$. Now γ_s has an endpoint at a , and we have $\gamma_s \cap B_R(s) \subset N_H(S)$. From the inclusions $S \cap B_R(s) \subset N_H(\gamma_s)$, and $\gamma_s \cap B_R(s) \subset N_H(S)$ and the inequality $6H < \frac{R}{4}$, we conclude that $N_H(\gamma_s) \cap S \subset S_0$. Choosing $x \in \gamma_s$ at distance $\frac{R}{2}$ from s , in the direction of a , we find some point $s' \in S_0$ with $d_\Gamma(x, s') \leq H$. See Figure 2 for a schematic. By repeated applications of the triangle inequality, one can easily show that this s' is further than $3H + 6\delta$ from γ , a contradiction.

We now argue that $S = S_0$. To see this, suppose that S_1 were some other component. We may apply the same argument to S_1 to produce a biinfinite geodesic γ_1 . The geodesics γ and γ_1 share an endpoint a , and so contain points within δ of one another. This implies that S_0 and S_1 contain points within $6H + 13\delta < \frac{R}{4}$ of one another, so they cannot be different $\frac{R}{4}$ -connected components.

We have established the first conclusion, since $d_{\text{Haus}}(S, S(a, \theta)) \leq 1$.

Now we suppose that $s \in S$ is in $\bar{\pi}(\sigma(D_0) \cap L_a \cap Y_\theta)$. By Lemma 6.8, we may take s to be in $\bar{\pi}((a, b, c), b)$ for some $b \in B_\epsilon(\theta)$. We can therefore take γ_s in the first part of this argument to be a geodesic joining a to b . The second conclusion of Lemma 2.12 implies that $(e^+ | b)_s \geq R - (4H + 10\delta)$, where e^+ is the endpoint of γ which is not equal to a . We thus have

$$\begin{aligned} (e^+ | b)_e &\geq (e^+ | b)_s - d_\Gamma(\mathbf{e}, s) \\ &\geq (e^+ | b)_s - \text{diam}(\pi(D_0)) \\ &\geq R - (4H + 10\delta). \end{aligned}$$

The first condition on ϵ in Lemma 6.3 implies that $(b | \theta)_e \geq R$. We have

$$\begin{aligned} (e^+ | \theta)_e &\geq \min\{(e^+ | b)_e, (b | \theta)_e\} - \delta \\ &\geq R - (4H + 11\delta), \end{aligned}$$

establishing the last claim of the proposition. □

7. The endpoint map

To summarize the results of the previous section, for each pair $a, \theta \in \partial G \times \partial G$ so that $L_a \cap Y_\theta \neq \emptyset$, there is a geodesic in Γ at (uniformly) bounded Hausdorff distance from $S(a, \theta) = \bar{\pi}(L_a \cap Y_\theta)$, with one endpoint equal to a . Say this geodesic is *shadowed* by $S(a, \theta)$, and orient it so that the negative endpoint is a . Any two biinfinite geodesics shadowed by $S(a, \theta)$ are bounded Hausdorff distance from each other, so they have the same endpoints in ∂G . This gives us positive and negative “endpoint maps” e^+ and e^- assigning to each pair (a, θ) , where $L_a \cap Y_\theta \neq \emptyset$ the positive and negative endpoints of the shadowed geodesic. For any such (a, θ) , we have $e^-(a, \theta) = a$ and $e^+(a, \theta) \neq a$. Furthermore, the equivariance property in Proposition 3.2 implies that for any $g \in G$, we have

$$gS(a, \theta) = S(ga, \rho(g)\theta) = \bar{\pi}(L_{ga} \cap Y_{\rho(g)\theta}).$$

This implies the following equivariance of the positive endpoint map

$$g \cdot e^+(a, \theta) = e^+(ga, \rho(g)\theta). \tag{12}$$

Here, and in what follows, we will frequently omit ρ_0 from the notation when it is clear that we are referring to the natural action of G on its boundary. Thus, we will write ga rather than $\rho_0(g)a$ when a is a boundary point. For $x \in X$, we also write gx for its image under the standard action of $g \in G$ on $X \subset \Xi$.

We will first establish continuity of the positive endpoint map on a large set, then use it to define a semiconjugacy.

7.1. Continuity

Proposition 7.1 (Continuity in a, θ over D_1). *Suppose $a_0, \theta_0 \in \partial G \times \partial G$ is such that there exists c with $(a_0, \theta_0, c) \in D_1$. Then the map $(a, \theta) \mapsto e^+(a, \theta)$ is continuous at (a_0, θ_0) .*

Proposition 7.1 will be a quick consequence of the following technical lemma.

Lemma 7.2. *Suppose $L_{a_0} \cap Y_{\theta_0}$ is nonempty. For any $r > 0$, there is a neighborhood $N = N(r)$ of $(a_0, \theta_0) \in \partial G \times \partial G$ so that if $(a, \theta) \in N$, then $B_r(\mathbf{e}) \cap S(a_0, \theta_0)$ lies in the $\text{diam}(\pi(D_1))$ -neighborhood of $S(a, \theta)$.*

The main idea behind the proof of Lemma 7.2 comes from the proof of Lemma 6.8, and the fact that the fixed point property used there is stable under small perturbations of the map $f_{a,c,\theta}$.

Proof of Lemma 7.2. If $B_r(\mathbf{e}) \cap S(a_0, \theta_0)$ is empty, there is nothing to show, so we suppose $B_r(\mathbf{e}) \cap S(a_0, \theta_0)$ is nonempty. Let $K \subset X$ be the closure of $\pi^{-1}(B_r(\mathbf{e}))$.

Recall from Notation 6.7 that $f_{a,c,\theta}$ denotes the map $x \mapsto f_\theta(a, x, c)$, and recall that $D_0 \subset D_1$ has the following properties

1. $X \subset GD_0$
2. For any $(a, b, c) \in D_0$, the set $B_{\epsilon_1}(a) \times B_{\epsilon_1}(b) \times B_{\epsilon_1}(c)$ is contained in D_1 (see Lemma 6.2.(4)).

Recall also that $\epsilon_1 > \epsilon_2 > 2\epsilon$.

We cover $\sigma(K) \cap (L_{a_0} \cap Y_{\theta_0})$ with translates of D_0 , as follows. Since K is compact, there are finitely many elements $g_1, g_2, \dots, g_k \in G$ so that

$$\sigma(K) \cap (L_{a_0} \cap Y_{\theta_0}) \subset \bigcup_{i=1}^k \sigma(g_i D_0).$$

By deleting elements from the list if necessary, we may assume

$$\sigma(g_i D_0) \cap (L_{a_0} \cap Y_{\theta_0}) \neq \emptyset$$

for each i .

Our next goal is to show that, for each i , the projection of the larger translate $g_i D_1$ to the Cayley graph contains a point of $S(a, \theta)$, provided that (a, θ) is chosen close enough to (a_0, θ_0) . Here, “close enough” depends on the set K and hence on the constant r .

Translating back to D_0 , for each i , we have $\sigma(D_0) \cap (L_{g_i^{-1}a_0} \cap Y_{\rho(g_i)^{-1}\theta_0}) \neq \emptyset$. Let $b_i, c_i \in \partial G$ be such that

$$((g_i^{-1}a_0, b_i, c_i), b_i) \in \sigma(D_0) \cap (L_{g_i^{-1}a_0} \cap Y_{\rho(g_i)^{-1}\theta_0}).$$

Because $(g_i^{-1}a_0, b_i, c_i) \in D_0 \subset D_2$, Definition 6.4, Item (4) implies that

$$b_i = f_{\rho(g_i)^{-1}\theta_0}(g_i^{-1}a_0, b_i, c_i) \in B_\epsilon(\rho(g_i)^{-1}\theta_0).$$

Hence, $d_{\text{vis}}(b_i, \rho(g_i)^{-1}\theta_0) < \epsilon$, and so $(g_i^{-1}a_0, \rho(g_i)^{-1}\theta_0, c_i) \in D_1$.

Let $p_i = \rho(g_i)^{-1}\theta_0$. By Item (2) of Definition 6.4, there is some $s c'_i \in B_\epsilon(c_i)$ so that $\{a_i\} \times B_{\epsilon_2}(p_i) \times \{c'_i\}$ lies in the interior of X . Furthermore, by Lemma 6.3, Item (ϵ_2), there is a closed contractible set $B_{p_i} \subset \partial G$ with $B_\epsilon(p_i) \subset B_{p_i} \subset B_{\epsilon_2}(p_i)$, and

$$f_{g_i^{-1}a_0, c'_i, p_i}(B_{\epsilon_2}(p_i)) \subset B_\epsilon(p_i),$$

so $f_{g_i^{-1}a_0, c'_i, p_i}$ has a fixed point in $B_\epsilon(p_i)$. The property of taking the compact set B_{p_i} into the open ball $B_\epsilon(p_i)$ is open (in the compact-open topology on continuous maps), so also holds for any map sufficiently close to $f_{g_i^{-1}a_0, c'_i, p_i}$, provided the map is defined on B_{p_i} . Recall that the domain of definition of $f_{x,y,z}$ is the set $\{w \mid (x, w, z) \in X\}$. Thus, if a function $f_{x,y,z}$ is defined on a set $\{x\} \times B_{\epsilon_2}(p_i) \times \{z\}$ contained in the interior of X , and $B_{\epsilon_2}(p_i) \supset B_{p_i}$, then for all sufficiently close x', y', z' , the function $f_{x',y',z'}$ will be defined on B_{p_i} as well. Additionally, the functions $f_{x,y,z}$ vary continuously in the arguments (x, y, z) . Thus, we may take a neighborhood N_i of (a_0, θ_0) , such that for each $(a, \theta) \in N_i$,

1. the map $f_{g_i^{-1}a, c'_i, \rho(g_i)^{-1}\theta}$ is defined on B_{p_i}
2. the map $f_{g_i^{-1}a, c'_i, \rho(g_i)^{-1}\theta}$ has a fixed point contained in $B_\epsilon(p_i)$; and
3. $d_{\text{vis}}(g_i^{-1}a, g_i^{-1}a_0) < \epsilon_1$.

Set $N = \bigcap_{i=1}^k N_i$. Thus, for any $(a, \theta) \in N$, each of the sets $L_{g_i^{-1}(a)} \cap Y_{\rho(g_i)^{-1}\theta}$ contains a point $\sigma(g_i^{-1}a, z_i, c'_i)$, where $z_i \in B_\epsilon(p_i)$. In particular, we have

$$\begin{aligned} d_{\text{vis}}(g_i^{-1}a, g_i^{-1}a_0) &< \epsilon_1 \text{ and} \\ d_{\text{vis}}(z_i, b_i) &\leq 2\epsilon < \epsilon_1. \end{aligned}$$

Since $(g_i^{-1}a_0, b_i, c_i) \in D_0$, this means that $(g_i^{-1}a, z_i, c'_i) \in D_1$. Multiplying on the left by g_i , we obtain $(a, g_i z_i, g_i c'_i) \in g_i D_1$ and $\sigma(a, g_i z_i, g_i c'_i) \in L_a \cap Y_\theta$, so the intersection

$$\sigma(g_i D_1) \cap (L_a \cap Y_\theta)$$

is nonempty for each of the elements g_i . Projecting to the Cayley graph, we have $B_r(\mathbf{e}) \cap S(a_0, \theta_0)$ contained in the $\text{diam}(\pi(D_1))$ -neighborhood of $S(a, \theta)$, which proves the lemma. \square

Proof of Proposition 7.1 from Lemma 7.2. Suppose $(a_0, \theta_0) \in \partial G \times \partial G$ is such that there exists c with $(a_0, \theta_0, c) \in D_1$. Then by Lemma 6.8, $L_{a_0} \cap Y_{\theta_0}$ contains a point of $\sigma(D_2)$, so is nonempty. Lemma 7.2 states that, given $r > 0$, there is a neighborhood N of (a_0, θ_0) so that if $(a, \theta) \in N$, then $B_r(\mathbf{e}) \cap S(a_0, \theta_0)$ lies in the $\text{diam}(\pi(D_1))$ -neighborhood of $S(a, \theta)$. By Proposition 6.11, both $S(a_0, \theta_0)$ and $S(a, \theta)$ are Hausdorff distance at most $3H + 6\delta + 1$ from some biinfinite geodesic, so these geodesics will $2(3H + 6\delta + 1) + \text{diam}(\pi(D_1))$ fellow-travel each other over a compact set, which can be taken as large as we wish by taking r large. This gives continuity. \square

We now establish continuity of a similar positive endpoint map defined everywhere on X . Since f is a homeomorphism and the sets $X \times \{\theta\}$ partition $X \times \partial G$, so do their images Y_θ . Likewise, the sets L_a give a partition of $\sigma(X) \subset X \times \partial G$, so for each $x \in X$, there exists a unique $a(x)$ and $\theta(x)$, such that $\sigma(x) \in L_{a(x)} \cap Y_{\theta(x)}$. Note if $x = (a, b, c)$, then $a(x) = a$. We have a sequence of maps:

$$x = (a, b, c) \mapsto ((a, b, c), b) \mapsto (a, \theta(x)) \mapsto e^+(a, \theta(x)). \tag{13}$$

Proposition 7.3 (Continuity on X). *Let $\mathcal{E}^+(x) = e^+(a(x), \theta(x))$ be the map given by the composition in (13). Then \mathcal{E}^+ is continuous on all of X .*

Proof. Equivariance of each map in the composition implies that we have the equivariance property

$$\mathcal{E}^+(gx) = \mathcal{E}^+(ga(x), \rho(g)\theta(x)).$$

It therefore suffices to check continuity of the composition above on the set $D_0 \cap X$ containing a fundamental domain for the action of G on X .

By definition, the section σ is continuous. For the second map in (13), note that $\theta(x)$ is simply projection onto the second coordinate of $f^{-1}(\sigma(x)) \in X \times \partial G$. Since f^{-1} is a homeomorphism of $X \times \partial G$, its projection $\theta(x)$ is continuous. Finally, Proposition 7.1 says that $(a, \theta(a, b, c)) \mapsto e^+(a, \theta)$ is continuous for $(a, b, c) \in D_1$. \square

Going forward, we will abuse notation and often think of \mathcal{E}^+ as a map defined on $\sigma(X)$, via the identification of X with $\sigma(X)$.

7.2. Proof of Theorem 1.1

Using the work above, we may now conclude the proof of our main theorem. First recall the statement.

Theorem 1.1 (Topological stability). *Let G be a hyperbolic group with sphere boundary. Then the action of G on ∂G is topologically stable. More precisely, given any neighborhood V of the identity in the space of continuous self-maps of S^n , there exists a neighborhood U of the standard boundary action in $\text{Hom}(G, \text{Homeo}(S^n))$, such that any representation in U has ρ_0 as a factor, with semiconjugacy contained in V .*

Our neighborhood $U = \mathcal{U}(V)$ was determined by our desired lower bound C_V on Gromov products when we set our conventions in Section 6.1. In this section, we show that $e^+(a, \theta)$ is (locally) a function only of θ , hence can be thought of as a map from ∂G to ∂G . We will then show that this map has the properties of the desired semiconjugacy between ρ_0 and ρ .




Figure 3. Paths of endpoints (red/solid line) on $\partial G = \partial \Gamma$ and associated near-geodesic sets in Γ (blue/dotted) with endpoints in the ϵ -balls about p and q .

Lemma 7.4 (e^+ is locally a function of θ). Let $\theta \in \partial G$, and let $\{a_t \mid t \in [0, 1]\}$ be a path in ∂G so that $e^+(a_t, \theta)$ is defined and continuous at all points. Then $e^+(a_t, \theta)$ is constant.

Proof. We argue by contradiction. Suppose we have such a path where e^+ is nonconstant. Truncating and reparameterizing, we may suppose that

1. $e^+(a_t, \theta)$ is not locally constant at $t = 0$,
2. for all $t \in [0, 1]$, $a_t \neq e^+(a_0, \theta)$, and
3. $e^+(a_1, \theta) \neq e^+(a_0, \theta)$.

The second item can be ensured by taking any sufficiently short path that is nonconstant at 0, since $e^+(a, \theta) \neq a$ holds for all a, θ .

Since G acts on ∂G as a uniform convergence action (see Proposition 2.2), the point $e^+(a_0, \theta)$ is a conical limit point, so there exists a sequence $\{g_i\} \subset G$ and points $p \neq q \in \partial G$, such that $g_i e^+(a_0, \theta) \rightarrow p$ and $g_i z \rightarrow q$ for all $z \in \partial G - \{e^+(a_0, \theta)\}$. Modifying the sequence $\{g_i\}$ if needed by postcomposing with some fixed $g \in G$, we may also assume there exists c with $(p, q, c) \in D_0$. In particular, this implies that $B_\epsilon(p) \cap B_\epsilon(q) = \emptyset$. Since ∂G is compact, we may assume by passing to a subsequence that $\rho(g_i)\theta$ converges to some point $\theta_\infty \in \partial G$.

For i large enough, $g_i \cdot e^+(a_1, \theta) \in B_{\epsilon/2}(q)$ and $g_i \cdot e^+(a_0, \theta) \in B_{\epsilon/2}(p)$. Thus, for each sufficiently large i , there exists t_i , such that $p_i := g_i \cdot e^+(a_{t_i}, \theta)$ is visual distance exactly ϵ from p . Since the arc a_t does not meet $e^+(a_0, \theta)$, we have $g_i a_{t_i} \in B_\epsilon(q)$ for all t and sufficiently large i . See Figure 3 for a schematic illustration. Consider the sets $g_i S(a_{t_i}, \theta) = S(g_i(a_{t_i}), \rho(g_i)(\theta))$. Recall by (12), we have

$$e^+(g_i a_{t_i}, \rho(g_i)\theta) = g_i \cdot e^+(a_{t_i}, \theta).$$

This implies that the geodesics shadowed by the sets $S(g_i(a_{t_i}), \rho(g_i)(\theta))$ all pass through some compact subset of the Cayley graph Γ , and so the sets themselves all meet some compact $K \subset X$.

For each i , fix a point $y_i \in L_{g_i a_{t_i}} \cap Y_{\rho(g_i)(\theta)} \cap K$. After passing to a further subsequence, the points y_i converge to some $y_\infty \in K$. Now $g_i a_{t_i} \rightarrow q$ and $\rho(g_i)(\theta) \rightarrow \theta_\infty$, so $y_\infty \in L_q \cap Y_{\theta_\infty} \cap K$. Using the notation from Proposition 7.3, continuity of positive endpoints implies that $\mathcal{E}^+(y_\infty) = \lim_{n \rightarrow \infty} \mathcal{E}^+(y_i)$ which is by construction some point at distance ϵ from p .

Now consider instead the constant sequence $t = 0$ instead of t_i . By the same reasoning, for i sufficiently large, $L_{g_i a_0} \cap Y_{\rho(g_i)(\theta)}$ will contain a point z_i in K . After passing to a subsequence, these converge to a point $z_\infty \in L_q \cap Y_{\theta_\infty} \cap K$. By continuity of \mathcal{E}^+ , we have $\mathcal{E}^+(z_\infty) = \lim_{n \rightarrow \infty} \mathcal{E}^+(z_i) = p$. Thus, we have

found two points, y_∞ and z_∞ , both in $Y_{\theta_\infty} \cap L_q$ with different positive endpoints $\mathcal{E}^+(z_\infty) \neq \mathcal{E}^+(y_\infty)$. This directly contradicts Proposition 6.11, and this contradiction concludes the proof. \square

Recall that Item (6) from Lemma 6.2 says that, if (a, θ, c) and (a', θ, c') lie in D_1 , then there exists a path a_t with $a_0 = a$ and $a_1 = a'$ and a point c'' , such that $(a_t, \theta, c'') \in D_1$ for all t . Proposition 7.1 says that the map $e^+(a_t, \theta)$ is therefore continuous at each point, and thus by Lemma 7.4, we conclude it is constant. In summary, we have the following.

Corollary 7.5. *If (a, θ, c) and (a', θ, c') lie in D_1 , then $e^+(a, \theta) = e^+(a', \theta)$.*

Definition 7.6. Define $h : S^{n-1} \rightarrow S^{n-1}$ by $h(\theta) = e^+(a, \theta)$, where a is any point, such that there exists c with $(a, \theta, c) \in D_0$.

Note that h is defined everywhere, and is continuous by the continuity of $e^+(a, \theta)$ given by Proposition 7.1. It remains to check that h satisfies the other properties of the semiconjugacy required to prove Theorem 1.1. The second point in Proposition 6.11 states that $(h(\theta) | \theta)_e \geq R - (4H + 11\delta)$. Our choice of R in Lemma 6.2 (★) implies that $(h(\theta) | \theta)_e \geq C_V$, satisfying (†) as desired, and showing that h lies in our chosen neighborhood V . This neighborhood contains only degree one maps, so h is surjective.

We now check equivariance. Let g be an element of the generating set $S \cup S^{-1}$ used in the definition of Γ , and let θ be given. Choose a so that $(a, \theta, c) \in D_0$. We have

$$\rho_0(g)h(\theta) = ge^+(a, \theta) = e^+(ga, \rho(g)\theta).$$

By definition of $D_{\frac{1}{2}}$, and our conditions on ρ , we have $(ga, g\theta, gc) \in D_{\frac{1}{2}}$. By definition of D_1 , and Item (3) of Definition 6.4 which defines the neighborhood $\mathcal{U}(V)$, we then have $(ga, \rho(g)\theta, gc) \in D_1$. Thus, by Corollary 7.5, $e^+(ga, \rho(g)\theta, gc) = e^+(a', \rho(g)\theta, c')$ for any choice of a' and c' , such that $(a', \rho(g)\theta, c') \in D_0$, thus giving

$$\rho_0(g)h(\theta) = h(\rho(g)\theta) \tag{14}$$

for all $\theta \in \partial G$. Since (14) holds for generators of G , it holds, inductively, for all elements $g \in G$. This shows that h is a semiconjugacy in the specified neighborhood of the identity map of S^n , completing the proof of the theorem.

Acknowledgements. K.M. was partially supported by NSF grant DMS 1844516 and a Sloan fellowship. J.M. was partially supported by Simons Collaboration Grant #524176.

Competing interest. The authors have no competing interest to declare.

References

- [AMN11] H. Abels, A. Manoussos and G. Noskov, ‘Proper actions and proper invariant metrics’, *J. Lond. Math. Soc.* (2) **83**(3) (2011), 619–636.
- [BH99] M. R. Bridson and A. Haefliger, ‘Metric spaces of non–positive curvature’, in *Grundlehren der Mathematischen Wissenschaften* vol. **319** (Springer–Verlag, Berlin, 1999).
- [BLW10] A. Bartels, W. Lück and S. Weinberger, ‘On hyperbolic groups with spheres as boundary’, *J. Differential Geom.* **86**(1) (2010), 1–16.
- [BM19] J. Bowden and K. Mann, ‘ C^0 stability of boundary actions and inequivalent Anosov flows’, Preprint, 2019, [arXiv:1909.02324](https://arxiv.org/abs/1909.02324). [math.DS].
- [Bow98] B. H. Bowditch, ‘A topological characterisation of hyperbolic groups’, *J. Amer. Math. Soc.* **11**(3) (1998), 643–667.
- [Cal61] E. Calabi, ‘On compact, Riemannian manifolds with constant curvature. I’, in *Proceedings of Symposia in Pure Mathematics* vol. **III** (American Mathematical Society, Providence, RI, 1961), 155–180.
- [CJ94] A. Casson and D. Jungreis, ‘Convergence groups and Seifert fibered 3-manifolds’, *Invent. Math.* **118**(3) (1994), 441–456.
- [CP93] M. Coornaert and A. Papadopoulos, ‘Symbolic dynamics and hyperbolic groups’, in *Lecture Notes in Mathematics* vol. 1539 (Springer–Verlag, Berlin, 1993).
- [EK71] R. D. Edwards and R. C. Kirby, ‘Deformations of spaces of imbeddings’, *Ann. of Math.* (2) **93** (1971), 63–88.

- [Gab92] D. Gabai, ‘Convergence groups are Fuchsian groups’, *Ann. of Math. (2)* **136**(3) (1992), 447–510.
- [GdlH90] E. Ghys and P. de la Harpe (eds.), ‘Sur les groupes hyperboliques d’après Mikhael Gromov’, in *Progress in Mathematics* vol. 83 (Birkhäuser Boston, Inc., Boston, MA, 1990). Papers from the Swiss Seminar on Hyperbolic Groups held in Bern, 1988.
- [Ghy01] E. Ghys, ‘Groups acting on the circle’, *Enseign. Math. (2)* **47**(3–4) (2001), 329–407.
- [GMRS98] R. Gitik, M. Mitra, E. Rips and M. Sageev, ‘Widths of subgroups’, *Trans. Amer. Math. Soc.* **350**(1) (1998), 321–329.
- [Gro87] M. Gromov, ‘Hyperbolic groups’, in Essays in group theory, *Mathematical Sciences Research Institute Publications* vol. 8 (Springer, New York, 1987), 75–263.
- [Kap] M. Kapovich, ‘Solution to MathOverflow’, post 139025, “Fixed points on boundary of hyperbolic group”. <https://mathoverflow.net/questions/139025/fixed-points-on-boundary-of-hyperbolic-group>
- [KB02] I. Kapovich and N. Benakli, ‘Boundaries of hyperbolic groups’, in Combinatorial and geometric group theory (New York, 2000/Hoboken, NJ, 2001) vol. 296 *Contemporary Mathematics* (American Mathematical Society, Providence, RI, 2002), 39–93.
- [KK00] M. Kapovich and B. Kleiner, ‘Hyperbolic groups with low-dimensional boundary’, *Ann. Sci. École Norm. Sup. (4)* **33**(5) (2000), 647–669.
- [KKL] M. Kapovich, S. Kim and J. Lee, ‘Structural stability of meandering-hyperbolic group actions’, Preprint, 2022, [arXiv:1904.06921](https://arxiv.org/abs/1904.06921). [math.GR].
- [Mat87] S. Matsumoto, ‘Some remarks on foliated S^1 bundles’, *Invent. Math.* **90**(2) (1987), 343–358.
- [Min05] A. Minasyan, ‘On residualizing homomorphisms preserving quasiconvexity’, *Comm. Algebra* **33**(7) (2005), 2423–2463.
- [Mos73] G. D. Mostow, ‘Strong rigidity of locally symmetric spaces’, in *Annals of Mathematics Studies* vol. 78 (Princeton University Press, Princeton, NJ, 1973), v+195.
- [New31] M. H. A. Newman, ‘A theorem on periodic transformations of spaces’, *Q. J. Math.* **2**(1) (1931), 1–8.
- [Sul85] D. Sullivan, ‘Quasiconformal homeomorphisms and dynamics. II. Structural stability implies hyperbolicity for Kleinian groups’, *Acta Math.* **155**(3–4) (1985), 243–260.
- [Tuk98] P. Tukia, ‘Conical limit points and uniform convergence groups’, *J. Reine Angew. Math.* **501** (1998), 71–98.
- [Wei62] A. Weil, ‘On discrete subgroups of Lie groups. II’, *Ann. of Math. (2)* **75** (1962), 578–602.