UNIMODULAR MEASURES ON THE SPACE OF ALL RIEMANNIAN MANIFOLDS

MIKLOS ABERT AND IAN BIRINGER

Abstract. We study unimodular measures on the space $\mathcal{M}^d$ of all pointed Riemannian d-manifolds. Examples can be constructed from finite volume manifolds, from measured foliations with Riemannian leaves, and from invariant random subgroups of Lie groups. Unimodularity is preserved under weak* limits, and under certain geometric constraints (e.g. bounded geometry) unimodular measures can be used to compactify sets of finite volume manifolds. One can then understand the geometry of manifolds M with large, finite volume by passing to unimodular limits.

We develop a structure theory for unimodular measures on $\mathcal{M}^d$, characterizing them via invariance under a certain geodesic flow, and showing that they correspond to transverse measures on a foliated 'desingularization' of $\mathcal{M}^d$. We also give a geometric proof of a compactness theorem for unimodular measures on the space of pointed manifolds with pinched negative curvature, and characterize unimodular measures supported on hyperbolic 3-manifolds with finitely generated fundamental group.

1. Introduction

The focus of this paper is on the class of 'unimodular' measures on the space

$$\mathcal{M}^d = \{\text{pointed Riemannian } d\text{-manifolds } (M,p)\}/\text{pointed isometry};$$

Throughout the paper, all Riemannian manifolds we consider are connected and complete. We consider $\mathcal{M}^d$ with the smooth topology, where two pointed manifolds (M,p) and (N,q) are smoothly close if there are compact subsets of M and N containing large radius neighborhoods of the base points that are diffeomorphic via a map that is C^1-close to an isometry, see §A.1.

Let $\mathcal{M}^d_2$ be the space of isometry classes of doubly pointed Riemannian d-manifolds (M,p,q), considered in the appropriate smooth topology, see §2.

Definition 1.1. A σ-finite Borel measure μ on $\mathcal{M}^d$ is unimodular if and only if for every nonnegative Borel function $f : \mathcal{M}^d_2 \to \mathbb{R}$ we have:

$$\int_{(M,p)\in\mathcal{M}^d}\int_{q\in M} f(M,p,q)\,dq\,d\mu = \int_{(M,p)\in\mathcal{M}^d}\int_{q\in M} f(M,q,p)\,dq\,d\mu. \tag{1}$$

Here, (1) is usually called the mass transport principle or MTP. One sometimes considers f to be a 'paying function', where $f(M,p,q)$ is the amount that the point (M,p) pays to (M,q), and the equation says that the expected income of a μ-random element $(M,p)\in\mathcal{M}^d$ is the same as the expected amount paid. Note that two sides of the mass transport principle can be considered as integrals $\int f\,d\mu_l$ and $\int f\,d\mu_r$ for two appropriate 'left' and 'right' Borel measures μ_l,μ_r on $\mathcal{M}^d_2$, so μ is unimodular if and only if $\mu_l=\mu_r$. See the beginning of §2.

Example 1.2 (Finite volume manifolds). Suppose M is a nite volume[1] Riemann-ian d-manifold, and let $_M$ be the measure on $\mathcal{M}^d$ obtained by pushing forward the Riemannian measure vol_M under the map

$$(2) \qquad M \to \mathcal{M}^d; \quad p \to (M; p):$$

Then both sides of the mass transport principle are equal to the integral of $f(M; p; q)$ over $(p; q) \in M \times M$, equipped with the product Riemannian measure, so the measure $_M$ is unimodular.

More generally, one can construct a nite unimodular measure on $\mathcal{M}^d$ from any Riemannian M that regularly covers a nite volume manifold. The point is that because of the symmetry given by the action of the deck group, the image of M in $\mathcal{M}^d$ actually looks like the nite volume quotient. See Example 2.4.

Example 1.3 (Measured foliations). Let X be a foliated space, a separable metriz-able space X that is a union of 'leaves' that t together locally as the horizontal factors in a product $\mathbb{R}^d \times Z$ for some transversal space Z. Suppose X is Riemann-ian, i.e. the leaves all have Riemannian metrics, and these metrics vary smoothly in the transverse direction. (See §3 for details.) There is then a Borel[2] leaf map $X \to \mathcal{M}^d; \; x \to (L_x; x)$; where L_x is the leaf through x.

Suppose that is a nite completely invariant measure on X, that is, a measure obtained by integrating the Riemannian measures on the leaves of X against some invariant transverse measure, see [35]. The push forward of under the leaf map is a unimodular measure on $\mathcal{M}^d$, see Theorem 1.8.

Example 1.4 (Many transitive manifolds). Let X be a Riemannian manifold with transitive isometry group, and note that any base point for X gives the same element $X \in \mathcal{M}^d$. In Proposition 2.6, we show that an atomic measure on $X \in \mathcal{M}^d$ is unimodular if and only if $\mathrm{Isom}(X)$ is a unimodular Lie group. Examples where $\mathrm{Isom}(X)$ is unimodular include nonpositively curved symmetric spaces, e.g. $\mathbb{R}^d; \mathbb{H}^d; SL_n\mathbb{R}/SO(n)$, and compact transitive manifolds like S^d. An example where $\mathrm{Isom}(X)$ is non-unimodular is the 3-dimensional Lie group $\mathrm{Sol}(p; q)$, where $p \neq q$, equipped with any left invariant metric[3].

Proposition 2.6 is one reason why these measures are called unimodular, al-though the mass transport principle also has a formal similarity to unimodularity of topological groups, being an equality of two 'left' and 'right' measures.

1.1. Motivation. Though this paper rst appeared online in 2016, we have rewrit-ten the section in 2020 to indicate how this paper ts into the eld currently. As such, many of the papers referenced below appeared after this one.

There are two main reasons to study unimodular measures. First, the space $\mathcal{M}^d$ provides a convenient universal setting in which to view nite volume manifolds, measured foliations, and innite volume manifolds that have a sucient amount

[1]If M has innite volume, the push forward measure $_M$ still satises the mass transport principle, but may not be -nite. For instance, if X has transitive isometry group then the map $X \to \mathcal{M}^d$ is a constant map, and $_X$ is only -nite if X has nite volume. On the other hand, if the isometry group of X is trivial, $_X$ will always be -nite.

[2]This is Proposition 4.1, where in fact we show that the leaf map is upper semi-continuous in a certain sense, extending a result of Lessa [68].

[3]In [28, Lemma 2.6], it is shown that the isometry group is a nite extension of $\mathrm{Sol}(p; q)$, which is not unimodular when $p \neq q$.

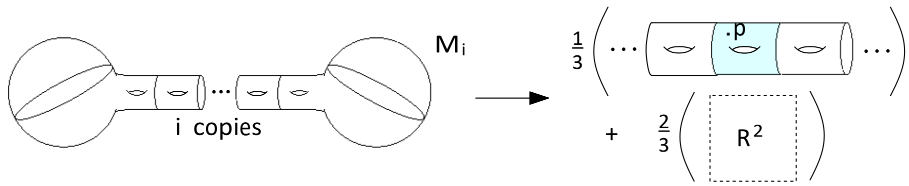

Figure 1. Create surfaces M_i by gluing i copies of some unit volume surface T end-to-end via a xed gluing map, capping o with two round spheres, each with volume i, and smoothing the result. Here, $_{M_i}$ =vol(M_i) weak* converges to a convex combination of an atomic measure on the single point R^2 in M^d, and a probability measure constructed by gluing bi-innitely many copies of T together and choosing a random base point from, say, the center copy.

of symmetry (e.g. the transitive examples above). More importantly, though, interpreting all these examples as measures on a single universal space allows one to dene a notion of convergence from one to another through weak* convergence of the associated measures on M^d. Here, recall that $_i$! in the weak* topology (or for some authors, the weak topology) if $\int f\, d_i \;!\; \int f\, d$ for every bounded, continuous function $f : M^d \;!\; R$.

An important special case of weak* convergence to keep in mind is when the measures $_i$ = $_{M_i}$ =vol(M_i), for some sequence of nite volume manifolds (M_i), as in Example 1.2. If $_i$! in the weak* topology, we say that the sequence (M_i) Benjamini{Schramm (BS) converges to , see e.g. [3]. This is in honor of an analogous notion introduced by those two authors in graph theory [18]. See also §1.2 below for more of this history. In some sense, the limit measure encodes, for large i, what the geometry of M_i looks like near randomly chosen base points, up to small metric distortion. As an example, consider Figure 1. The transition between the spheres and the neck is lost in the weak* limit, since the probability that a randomly chosen base point will lie near there is negligible, and the small metric distortion allows R^2 to approximate the large radius spheres.

As mentioned in §2, unimodularity is preserved under weak* limits, so in particular, all BS-limits of nite volume manifolds are unimodular[4]. So, if one can develop a robust theory of unimodular measures, one can then try to use this theory to analyze sequences of nite volume manifolds via their weak* limits, as above.

As an example, let X be a irreducible symmetric space of noncompact type. An X-manifold is a quotient M = nX, where acts freely, properly discontinuously and isometrically. In recent joint work with Bergeron and Gelander [1], using the framework developed in this paper we showed:

Theorem 1.5 (Abert, Bergeron, Biringer, Gelander [1]). Suppose that X is not a metric scale of H^3. If (M_i) is any BS-convergent sequence of nite volume X-manifolds, then for all k the sequence $b_k(M_i)$=vol(M_i) converges.

Essentially, this means that in the context above, a given limit measure has some sort of 'Betti number' that is the limit of the volume normalized Betti numbers of any approximating sequence of M_i's. It would be interesting to develop an

[4]The converse is open, and is essentially equivalent to the analogous question for unimodular random graphs, which generalizes the open question of whether all groups are soc, see [9].

intrinsic denition of such a Betti number for an arbitrary unimodular measure on M^d. Relatedly, an analogous denition has been recently given by Michael Schrödl for unimodular random rooted simplicial complexes, see [87]. However, the manifold case presents additional diculties.

One situation in which there is an existing denition, though, is when the unimodular measure is just an atomic measure $_X$ on the single point X 2 M^d. In this case, the appropriate invariants are the L^2-Betti numbers $^{(2)}_k(X)$, see e.g. [3] for denitions. In our 2012 paper with Abert et al [3], we all had previously shown that if $\mathrm{rank}_{\mathbb{R}}(X)$ 2, then any sequence of distinct nite volume X-manifolds BS-converges to $_X$. Combining this with Theorem 1.5 above gives:

Theorem 1.6 (Corollary 1.4 of [1]). Suppose that $\mathrm{rank}_{\mathbb{R}}(X)$ 2 and (M_i) is any sequence of distinct nite volume X-manifolds. Then for all k 2 N, we have

$$b_k(M_n)/\mathrm{vol}(M_n) \; ! \; {}^{(2)}_k(X):$$

This extends earlier work with Nikolov-Raimbault-Samet in [3], and is a uniform version of Lück's approximation theorem [74] that applies to all quotients of a xed symmetric space, not just covers of a single quotient. We stress that it is necessary in [1] to work geometrically, and so most of that paper is written using the framework of convergence of measures on M^d, as developed in this paper.

Subsequent work of Abert{Bergeron{Masson [4] exploits the language of Benjamini-Schramm convergence of manifolds introduced in this paper to analyze eigenfunctions of the Laplacian for compact Riemannian manifolds. The asymptotic behavior of eigenfunctions has been studied extensively in the literature. There are two major directions of interest: one can study the eigenvalue aspect, where one has a xed manifold and the energy of the eigenfunction tends to innity, and the level aspect, where one looks at a covering tower of manifolds (usually coming from a subgroup chain for an arithmetic lattice) and the energy converges to a xed value. It was understood in the community that these aspects are related and most theo-rems on one side tend to nd their counterparts on the other side. Our language of Benjamini-Schramm convergence now unies the level and eigenvalue aspects. In-deed, for a covering tower, the limit of eigenfunctions will be an invariant random eigenwave on the limiting space (usually the symmetric space of the corresponding Lie group). For a xed manifold, rescaling the Riemannian metric with the energy will produce a sequence of manifolds and a xed eigenvalue, hence the limiting eigenwave will live on the standard Euclidean space. The language of Benjamini-Schramm convergence, in particular, allows one to give the rst mathematically precise formulation for the famous Berry conjecture in physics, and connects the conjecture to Quantum Unique Ergodicity. See [4] for details.

1.2. History and related papers. Much of our work is inspired by a recent program in graph theory, in particular work of Aldous-Lyons [9] and Benjamini-Schramm [18]. For instance, the term 'unimodularity' was previously used in [9], for measures on the space

$$G = \text{ rooted, connected, locally nite graphs } (G;p) \; = \text{automorphism}$$

such that for any Borel function f on the space of doubly rooted graphs,

$$(3) \qquad \int_{(G;p)2G} \sum_{q2V(G)} f(G;p;q)\,dq\,d = \int_{(G;p)2G} \sum_{q2V(G)} f(G;q;p)\,dq\,d:$$

In fact, this version of the mass transport principle appeared even earlier in [18], generalizing a concept important in percolation theory [16, 57].

As in Example 1.2, every nite graph G gives a unimodular measure $_G$ on G, by pushing forward the counting measure on its vertices under the map

$$G \quad ! \quad G; \quad p \quad ! \quad (G; p):$$

Similarly, a transitive graph gives a unimodular measure on G if and only if its automorphism group is unimodular, see [16] and [75, Section 8.2]. One can study unimodular measures on G that are weak* limits of the $_G$, c.f.[6, 18, 80], and extending results known for nite graphs to arbitrary unimodular measures on G has recently become a small industry, see e.g. [9, 16, 17, 57].

Ideas similar to ours have also appeared previously in the continuous setting, even apart from ABBGNRS [3]. Most directly, Bowen [25] used unimodular measures on the space of pointed metric measure spaces to bound the Cheeger constants of hyperbolic 4-manifolds with free fundamental group. In his thesis, Lessa [68], see also [69, 70], studied measures on M^d that are stationary under Brownian motion, of which unimodular measures are examples[5], and a few of the technical parts of this paper are similar to parts of his. Namazi-Panka-Souto [80], analyzed weak*-limits of the measures $_{M_i}$=vol(M_i) for sequences (M_i) of manifolds that are all quasi-conformal deformations of a xed closed manifold and that all have bounded geometry. Also, Vadim Kaimanovich has for some time promoted measured foliations from a viewpoint similar to ours, and we refer the reader to his papers [64, 63, 65] for culture.

1.3. Statements of results.

Most of the paper concerns the structure theory of unimodular measures. The case where is a unimodular probability measure is of particular interest: a -random element of M^d is then called a unimodular random manifold (URM). In this section, we will start by explaining the close relationship between unimodular measures and completely invariant measures on foliated spaces, as mentioned in Example 1.3. Then, we will outline the dictionary between invariant random subgroups and unimodular random locally symmetric spaces. As an interesting trip to the zoo, a characterization of unimodular random hyperbolic 2 and 3-manifolds with nitely generated fundamental group is given. We then discuss conditions under which sets of unimodular measures on M^d are weak* compact, and nish with a discussion of the rather long appendix, where it is shown that M^d and various related spaces have reasonable topology.

1.3.1. Unimodularity and foliated spaces.

As mentioned above, a separable, metrizable space X is a foliated space if it is a union of leaves that t together locally as the horizontal factors in a product R^d Z for some transversal space Z. We say X is Riemannian if the leaves all have Riemannian metrics, and if these metrics vary smoothly in the transverse direction. See §3 for details.

[5] On foliated spaces, Lessa's 'stationary measures' correspond to harmonic measures, while our unimodular measures correspond to completely invariant measures. See §3, [35] and [68]. Also, see Benjamini{Curien [15] for a corresponding theory of stationary random graphs.

On such an X, let L_p be the leaf through p. A σ-finite Borel measure μ on X is unimodular if for every nonnegative Borel $f : X \times X \to \mathbb{R}$ we have:

$$\int_{p\in X}\int_{q\in L_p} f(p,q)\, d\mathrm{vol}_{L_p}\, d\mu = \int_{p\in X}\int_{q\in L_p} f(q,p)\, d\mathrm{vol}_{L_p}\, d\mu. \tag{4}$$

Also, a measure μ on X is called completely invariant if it is obtained by integrating the Riemannian measures on the leaves of X against some invariant transverse measure, see §3 and [35]. We then have:

Theorem 1.7. Suppose that X is a Riemannian foliated space and μ is a σ-finite Borel measure on X. Then the following are equivalent:

1) μ is completely invariant,
2) μ is unimodular,
3) μ lifts uniformly to a measure $\tilde\mu$ on the leaf-wise unit tangent bundle T^1X that is invariant under leaf-wise geodesic flow.

To understand the 'uniform lift' in 3), take a measure μ on X, and integrate against the (round) Riemannian measures on all leaf-wise tangent spheres T^1_pL to get a measure $\tilde\mu$ on T^1X. A version of this condition will reappear below as an alternative characterization of unimodularity for measures on $\mathcal{M}^d$: Theorem 1.7 is proved in §3, where it is restated as Theorem 3.1. Two additional characterizations of unimodularity are included in the new statement, one of which parallels a well-known result in graph theory.

In some sense, the space $\mathcal{M}^d$ is itself almost foliated, where the 'leaves' are the subsets obtained by fixing a manifold M and varying the basepoint $p \in M$. One would like to say that unimodular measures are just completely invariant measures, with respect to this foliation. However, due to the equivalence relation defining $\mathcal{M}^d$, these 'leaves' are actually of the form $M/\mathrm{Isom}(M)$, so the foliation is highly singular, and complete invariance does not make sense.

However, there is a way to make this point of view precise. Recall that if X is any Riemannian foliated space, its leaf map takes $x \in X$ to the pointed Riemannian manifold (L_x, x), where L_x is the leaf through x. We then have:

Theorem 1.8 (Desingularizing $\mathcal{M}^d$). If μ is a completely invariant probability measure on a Riemannian foliated space X, then μ pushes forward under the leaf map to a unimodular probability measure on $\mathcal{M}^d$.

Conversely, there is a Polish Riemannian foliated space $\mathcal{P}^d$ such that any σ-finite unimodular measure on $\mathcal{M}^d$ is the push forward under the leaf map of some completely invariant measure on $\mathcal{P}^d$. Moreover, for any fixed manifold M, the preimage of $\{(M,p) \mid p \in M\} \subset \mathcal{M}^d$ under the leaf map is a union of leaves of $\mathcal{P}^d$, each of which is isometric to M.

This theorem indicates an advantage our continuous framework has over graph theory: although the mass transport principle (3) does indicate a compatibility between unimodular measures on $\mathcal{G}$ and the counting measures on the vertex sets of fixed graphs G, there is no precise statement saying that unimodular measures are made by locally integrating up these counting measures in analogy with Theorem 1.8. In some sense, the problem is that graphs do not have enough local structure for this perspective to translate.

Alvarez Lopez and Barral Lijo [72] independently prove a desingularization theorem similar to Theorem 1.8, which they use to show that any manifold with bounded geometry can be realized isometrically as a leaf in a compact Riemannian foliated space. As their goals are topological, rather than measure theoretic, their foliated space is not set up so that one can lift measures on M^d to the foliated space using Poisson processes, though, a property that is crucial in our applications.

The idea behind the construction of P^d is simple. Since the problem is that Riemannian manifolds may have nontrivial isometries, we set P^d to be the set of isometry classes of triples $(M; p; D)$ where $p \in M$ is a base point and $D \subset M$ is a closed subset such that there is no isometry $f : M \to M$ with $f(D) = D$. The leaves are obtained by xing M and D and varying p, and the leaf map is just the projection $(M; p; D) \to (M; p)$. However, it takes some work to see that these leaves t together locally into a product structure $R^d \times Z$: a brief sketch of this argument is given in the beginning of §4.3. Assuming this, though, measures on M^d induce measures on P^d after integrating against a Poisson process on each ber of the leaf map. See §4.2 for details.

Completely invariant measures on foliated spaces have been well studied, e.g. [33, 46, 52, 51, 53]. So for instance, one can now take a sequence of nite volume manifolds (M_i), pass to the associated unimodular probability measures $_{M_i}/\mathrm{vol}(M_i)$, extract a weak* limit measure , study this using tools from foliations, and deduce results about the manifolds M_i.

For those working in foliations, the mass transport principle (1) may seem less interesting now that we know unimodularity can also be characterized in terms of complete invariance. However, we would like to stress that often, the MTP is the more convenient denition to use. We illustrate this in Theorem 1.12, where we use the MTP to give a proof of weak* compactness of the set of unimodular measures supported on manifolds with pinched negative curvature. For another example, Biringer-Raimbault [19] have studied the space of ends of a unimodular random manifold, showing for example that it has either $0; 1$ or 2 elements, or is a Cantor set. This parallels a result of Ghys [53] on the topology of generic leaves of a measured foliation. Neither of these results quite implies the other, although Ghys's result is really more general, as it applies to harmonic measures, and not just completely invariant ones. However, the MTP encapsulates a recurrence that makes the proof in [19] extremely short.

One other reason to prefer unimodularity in our setting is that to talk about complete invariance, one must leave M^d, passing to an associated foliated space using the desingularization theorem. On the other hand, the geodesic ow invariance of Theorem 1.7 can be phrased (more or less) directly within M^d.

Theorem 1.9. Suppose that is a Borel measure on M^d. Then is unimodular if and only if its uniform lift ~ on T^1M^d is geodesic ow invariant.

See §4.2 for the proof. Here, T^1M^d is the space of isometry classes of rooted unit tangent bundles $(T^1M; p; v)$, where $v \in T^1_pM$. Each ber T^1_pM of

$$T^1M^d \to M^d; \quad (M; p; v) \to (M; p)$$

comes with a natural Riemannian metric induced by the inner product on T_pM, and we write $\omega_{M;p}$ for the associated Riemannian measure on T^1_pM. Then ~ is the measure on T^1M^d dened by the equation $d\tilde{} = \omega_{M;p}\, d$: The geodesic ows on individual T^1M combine to give a continuous ow

$$g_t : T^1M^d \to T^1M^d$$

and this is the geodesic flow referenced in the statement of the theorem.

This theorem is an analogue of a result in graph theory. Let $G_d \subset G$ be the space of isometry classes of pointed d-regular graphs. The associated space E_d of d-regular graphs with a distinguished oriented edge projects onto G_d, where the map replaces a distinguished edge with its original vertex. For each $(G; v) \in G_d$, the uniform probability measure on the set of d edges originating at v quotients to a probability measure on the fiber over $(G; v)$ in E_d. Integrating these fiber measures against gives a measure ~ on E, and Aldous-Lyons [9] proved that is unimodular if and only if ~ is invariant under the map $E \to E$ that switches the orientation of the distinguished edge.

1.3.2. Unimodular random manifolds and IRSs. We now focus on unimodular probability measures on M^d, in which case a -random element of M^d is called a unimodular random manifold (URM). There is a close relationship between URMs and invariant random subgroups (IRSs), which have been studied in [21, 24, 26, 44, 54, 59, 58, 89].

Let G be a locally compact, second countable group, and let Sub_G be the space of closed subgroups of G, endowed with its Chabauty topology, see A.4.

Definition 1.10. An invariant random subgroup (IRS) of G is a random element of a Borel probability measure on Sub_G that is invariant under the conjugation action of G on Sub_G.

When G is finitely generated, say by a symmetric set S, there is a dictionary between IRSs of G and unimodular random S-labeled graphs, or URSGs, which we will briefly explain. An S-labeled graph is a countable directed graph with edges labeled by elements of S, such that the edges coming out from any given vertex v have labels in 1-1 correspondence with elements of S, and the same is true for the labels of edges coming into v. Every subgroup $H < G$ determines an S-labeled Schreier graph $\mathrm{Sch}_S(H\backslash G)$, whose vertices are right cosets Hg and where each $s \in S$ contributes a labeled edge from every coset Hg to Hgs. Note that $\mathrm{Sch}_S(H\backslash G)$ comes with a natural base point, the identity coset H.

In a variation of the discussion in §1.2, let G_S be the space of isomorphism classes of rooted S-labeled graphs. A URSG is a random element of G_S with respect to a probability measure that satisfies the appropriate S-labeled analogue of the mass transport principle (3). A random subgroup $H < G$ determines a random rooted Schreier graph, and conjugation invariance of the distribution of H is equivalent to unimodularity of $\mathrm{Sch}_S(H\backslash G)$. So, IRSs of G exactly correspond to URSGs. See [5] and [21, §4] for details.

In the continuous setting, IRSs were first studied in ABBGNRS [3]. In analogy with the above, when a group G acts on X by isometries, there should be a dictionary between IRSs of G and certain unimodular random X-manifolds. Here, an X-manifold is just a quotient $\backslash X$, where acts freely and properly discontinuously on X by isometries. Two parts of this dictionary are discussed in §2.8, in the cases where G acts transitively, or discretely on X. The following is a particularly nice case of our analysis of IRSs of transitive G.

Proposition 1.11 (URMs vs IRSs). Suppose X is a simply connected Riemannian manifold whose isometry group is unimodular and acts transitively. Then there is a weak*-homeomorphism between the spaces of distributions of discrete, torsion free IRSs of Isom(X) and of unimodular random X-manifolds.

So, X could be a non-positively curved symmetric space, for instance $\mathbb{R}^d$, $\mathbb{H}^d$ or $SL_n\mathbb{R}/SO(n)$. Note that when we say an IRS or URM has a property like 'torsion free' or 'X', this property is to be assumed to be satised almost always. For instance, the above says that there is a homeomorphism between the space of conjugation invariant probability measures on Sub_G such that -a.e. $H \in \mathrm{Sub}_G$ is discrete and torsion free, and the space of unimodular probability measures on $\mathcal{M}^d$ such that -almost every (M, p) is an X-manifold.

1.3.3. Compactness theorems. To understand sequences of nite volume manifolds, then, one would naturally like to understand conditions under which sets of unimodular probability measures are compact, so that a unimodular weak* limit of the measures $_{M_i}/\mathrm{vol}(M_i)$ can be extracted after passing to a subsequence.

By work of Cheeger and Gromov, c.f. [56] and [81, Chapter 10], the subset of $\mathcal{M}^d$ consisting of pointed manifolds (M, p) with bounded geometry is compact. Here, bounded geometry means that the sectional curvatures of M, and all of their derivatives, are uniformly bounded above and below, and the injectivity radius at the base point p is bounded away from zero. See §5. By the Riesz representation theorem and Alaoglu's theorem, this implies that the set of unimodular probability measures supported on manifolds with bounded geometry is weak* compact, since unimodularity is a weak* closed condition.

Both the curvature bounds and the bound on injectivity radius are necessary for compactness of pointed manifolds. However, we show that in the presence of pinched negative curvature, an injectivity radius bound is unnecessary for compactness once we pass to measures:

Theorem 1.12. The set of all unimodular probability measures on $\mathcal{M}^d$ that are concentrated on pointed manifolds with pinched negative curvature and uniform upper and lower bounds on all derivatives of curvature is weak* compact.

See §5 for a more precise statement. The condition on the derivatives of curvature is only necessary because we consider $\mathcal{M}^d$ with the smooth topology; a topology of weaker regularity would require weaker assumptions. Essentially, the reason the injectivity radius assumption is not necessary is because in pinched negative curvature, the -thin part of a manifold takes up at most some uniform proportion $C()$ of the total volume, where $C \to 0$ as $\to 0$. In fact, Theorem 1.12 boils down to a precise version of this kind of statement, see the proof of Proposition 5.1, that still applies to manifolds with innite volume.

We explain in §5 that there is no analogue of Theorem 1.12 in nonpositive curvature, but using work from ABBGNRS [3], one can show that we still have a weak* compactness theorem for locally symmetric spaces:

Theorem 1.13. Let X be a symmetric space of nonpositive curvature with no Euclidean factors, and let $\mathcal{M}^X \subset \mathcal{M}^d$ be the subset of pointed X-manifolds. Then the space of unimodular probability measures on $\mathcal{M}^X$ is weak*-compact.

The proof of Theorem 1.13 is algebraic: it uses the dictionary between unimodular measures and IRSs discussed in the previous section, and arguments related to Borel's density theorem, c.f. [49]. We give this proof in §5.2, and also briey discuss the question of whether there is a universal theorem that generalizes both Theorems 1.12 and 1.13.

1.3.4. Hyperbolic 3-manifolds with nitely generated fundamental group, and The No-Core Principle. Finite volume hyperbolic manifolds have nitely generated $_1$, [14]. While the converse is not true in general, the question is at least interesting for d-manifolds M with enough symmetry: is it true that when M = H^d regularly covers a nite volume d-manifold and $_1$M is nitely generated, then M has nite volume?

When d = 2, the answer is yes. Any surface S with nitely generated fundamental group is geometrically nite, see [67, Theorem 4.6.1]. If S = H^2 regularly covers a nite volume surface, its limit set is the entire circle @$_1H^2$, say by [83, Theorem 12.2.14], and then geometric niteness implies that S has nite volume, see [67, Theorem 4.5.1]. The question is open for d 4.

When d = 3, Thurston's bered hyperbolization theorem [91] states that the mapping torus M of a pseudo-Anosov homeomorphism of a surface S admits a hyperbolic metric. The fundamental group of M splits as the semidirect product

(5) 1 ! $_1$S ! $_1$M ! Z ! 1;

and the regular cover $\hat{M}$ corresponding to $_1$S is a hyperbolic 3-manifold with nitely generated fundamental group. However, it is a well-known consequence of the Tameness Theorem of Agol [7] and Calegari-Gabai [29] and Canary's covering theorem [32] that these $\hat{M}$ are the only examples when d = 3.

A unimodular random hyperbolic manifold (URHM) is, as should be expected, a random element of M^d with respect to a unimodular probability measure concentrated on pointed hyperbolic d-manifolds. Simple examples include a nite volume hyperbolic manifold with a randomly chosen base point, and the hyperbolic space H^d. (See Examples 1.2 and 1.4.)

Any regular cover of a hyperbolic manifold can be considered as a URHM, via Example 2.4. It turns out that URHMs have enough symmetry that the rigidity results for regular covers discussed above have analogues for URHMs with nitely generated fundamental group. For instance, it follows from [3, Proposition 11.3] that the limit set of a URHM M, with M = H^d, is always the entire boundary sphere @$_1H^d$. When d = 2, this means that any URHM with nitely generated $_1$ has nite volume, via the same argument as above.

When d = 3, we constructed examples in [3, §12.5] of IRSs (hence URHMs, by Proposition 1.11) with nitely generated $_1$ that are not regular covers of nite volume manifolds. However, these examples all have the same coarse geometric structure as the $\hat{M}$ examples above: they are all doubly degenerate hyperbolic 3-manifolds homeomorphic to S R, for some nite type surface S. See §6 for denitions. Here, we show that these are the only examples:

Theorem 1.14. Every unimodular random hyperbolic 3-manifold with nitely generated fundamental group either is isometric to H^3, has nite volume, or is a doubly degenerate hyperbolic structure on S R for some nite type S.

Here is another informal way to motivate Theorem 1.14. Suppose that M is a hyperbolic 3-manifold with nite, but large, volume. Randomly choose a point p 2 M and consider a neighborhood U 3 p with some xed radius R, which is large, but say not as large as vol(M). What can U look like geometrically? On the one hand, it could be a large embedded ball from H^3, while at the other extreme, it could have very complicated topology, requiring many elements to generate $_1$U. If

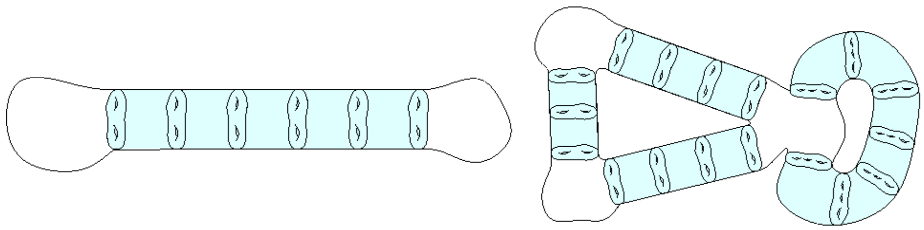

Figure 2. A schematic picturing a number of long product regions in hyperbolic 3-manifolds. Each is homeomorphic to a product S [0; 1], for some nite type surface S, and has bounded area level surfaces. If the product re-gions lengthen while the complexity of the underlying graph stays bounded, the associated probability measures $_{M_i}$ =vol(M_i) on M^d weak* converge to a URHM that is a doubly degenerate hyperbolic structure on S R.

$_1U$ can be generated by few elements, though, the geometry of U is more limited: essentially, it will look like a large piece of an innite volume hyperbolic 3-manifold N with some small number of ends.

Now, in any suciently large piece an innite volume N, the ends of N take up a much larger proportion of volume than the 'core' of N does. So, the probability that our base point p was randomly chosen to lie inside the core is negligible. In other words, most choices of p that end up in a neighborhood U with not so complicated topology will look like they are stuck deep inside an end of an innite volume hyperbolic 3-manifold. By the geometric classication of ends of hyperbolic 3-manifolds, c.f. [31] and also [32, The Filling Theorem], this means that the point p will either be contained in a large embedded ball from H^3 (when the end is geometrically nite) or a long product region (when the end is degenerate). See Figure 2.

Informally, this discussion means that near a randomly chosen point in a hyperbolic 3-manifold M with large nite volume, M will look like either

1) a large embedded ball from H^3,
2) a region for which the minimal number of $_1$-generators is 'very large',
3) a long product region.

To relate this back to Theorem 1.14, note that large hyperbolic balls and long product regions are exactly what one obtains by taking large neighborhoods within H^3 and doubly degenerate hyperbolic structures on S R, which are the only innite volume manifolds allowed in the theorem. And if a sequence (M_i) of nite volume manifolds with vol(M_i) ! 1 gives a sequence of probability measures $_{M_i}$ =vol(M_i) that weak* converges to a unimodular on M^d, the informal local analysis of the geometry of M_i described above translates exactly into Theorem 1.14. In fact, in light of the weak* compactness of the set of unimodular probability measures supported on hyperbolic manifolds, which follows from Theorem 1.12 or Theorem 1.13, one can view Theorem 1.14 as a precise version of the informal statement that for any large-volume M, the local geometry near a randomly chosen base point is as described above.

The key idea in the proof of Theorem 1.14 is the following:

Theorem 1.15 (The No-Core Principle). Suppose that μ is a unimodular probability measure on $\mathcal{M}^d$ and that $f : \mathcal{M}^d \to \{0,1\}$ is a Borel function. Then for μ-almost every element $(M,p) \in \mathcal{M}^d$, we have

$$0 < \mathrm{vol}_M\{q \in M \mid f(M,q) = 1\} < \infty \implies \mathrm{vol}(M) < \infty.$$

Geometrically, one should imagine that $f(M,p) = 1$ when the base point p lies in a 'core' of M. The theorem then says that when (M,p) lies in the support of a unimodular probability measure, one can only Borel-select a core with nite, nonzero volume for M when M has nite volume. While the statement above is very useful—it also is used in Biringer-Raimbault [19]—it is basically an immediate consequence of the mass transport principle, see §2.1.

Essentially, the proof of Theorem 1.14 is that any hyperbolic 3-manifold with nitely generated fundamental group has a nite volume core, obtained by chopping o neighborhoods of its innite volume ends. However, it requires some work to choose the core in a canonical enough way so that the function f in the No-Core Principle is Borel. See §6 for details.

1.3.5. Appendix: the topology of $\mathcal{M}^d$ and the Chabauty topology. The paper ends with a lengthy appendix. We give in §A.1 and A.6 slight extensions of existing compactness and stability results concerning smooth convergence, but most of the appendix is spent showing $\mathcal{M}^d$ and various related spaces are Polish. This is necessary to justify the use of measure theoretic tools like Rohlin's disintegration theorem, or Varadarajan's compact model theorem (see the proof of Proposition 4.18).

Candel, Alvarez López and Barral Lijo[10] have independently and concurrently studied the space $\mathcal{M}^d$, proving that it is Polish and establishing a number of interesting topological properties that are very related to this paper, e.g. to Proposition 4.1. Their proof was made publicly available before ours, so the result is really theirs. The two approaches are quite similar, but our proof produces an explicit metric that we use elsewhere in the paper, and is simpler in some ways, so we still present it here.

We nd the proofs that these spaces are Polish quite interesting. For instance, recall that two points (M,p) and (N,q) in $\mathcal{M}^d$ are smoothly close if there is a dieomorphism f from a large neighborhood $B \ni p$ to a large neighborhood of q, such that the metric $\langle,\rangle_N$ on N pulls back to a metric on B that is C^∞ close to $\langle,\rangle_M$, see §A.1. To metrize this denition directly, one would have to metrize the C^∞ topology on the appropriate space of tensors on M separately for each $(M,p) \in \mathcal{M}^d$, and then hope that the choice is canonical enough that the triangle inequality holds for the induced metric on $\mathcal{M}^d$. This is hard to do, so instead we dene the distance between (M,p) and (N,q) by measuring the bilipschitz distortion of the 'iterated total derivatives'

$$D^k f : T^k M \to T^k N$$

on the k-fold iterated tangent bundles $T^k M = T(T(T(M)))$, which we consider with the associated 'iterated Sasaki metrics'. See §A.2.

1.4. Plan of the paper. Section 2 introduces unimodular measures in detail, and discusses the No-Core Principle and the dictionary between URMS and IRSs. In Section 3, we review completely invariant measures on foliated spaces, and prove Theorem 3.1, which shows that complete invariance is equivalent to unimodularity, among other things. Section 4 discusses the foliated structure of $\mathcal{M}^d$, the desingularization theorem, and the ergodic decomposition of unimodular measures. The weak* compactness theorems mentioned in the introduction are proved in Section 5, while the characterization of unimodular random hyperbolic 3-manifolds with finitely generated π_1 is the main focus of Section 6. The paper ends with an appendix concerning the topology of $\mathcal{M}^d$ and related spaces.

1.5. Acknowledgments. We would like to thank Nir Avni, Igor Belegradek, Lewis Bowen, Renato Feres, Etienne Ghys, J. A. Alvarez Lopez, Juan Souto, Ralf Spatzier and Shmuel Weinberger for a number of useful discussions. We also thank a referee for very helpful comments, and for finding a significant error in the leadup to Theorem A.10. The second author was partially supported by NSF grant DMS-1308678.

2. Unimodular measures on $\mathcal{M}^d$

A rooted Riemannian d-manifold is a pair (M, p) where M is a Riemannian d-manifold and $p \in M$ is a basepoint. We assume that all Riemannian manifolds in this paper are complete and connected. A doubly rooted manifold is a triple (M, p, q) where M is a Riemannian d-manifold and $p, q \in M$.

Definition 2.1. Let $\mathcal{M}^d$ and $\mathcal{M}^d_2$ be the spaces of isometry classes of rooted and doubly rooted Riemannian d-manifolds, endowed with the smooth topology.

Recall that a sequence of rooted Riemannian manifolds (M_i, p_i) smoothly converges to (M, p) if for every $R > 0$, there is a C^∞-embedding

$$f_i : B_M(p, R) \to M_i$$

such that $f_i(p) = p_i$ and $f_i^*(g_{M_i}) \to g_M$ in the C^∞ topology, where g_{M_i} and g_M are the associated Riemannian metrics. The convergence $(M_i, p_i, q_i) \to (M, p, q)$ of a sequence of doubly rooted manifolds is the same, except that we require that $f_i(q) = q_i$ when defined. In §A.2, we show that smooth convergence comes from a Polish topology on $\mathcal{M}^d$. An analogous statement holds for $\mathcal{M}^d_2$.

Example 2.2. Setting $d = 1$, $\mathcal{M}^1$ is homeomorphic to $(0, \infty]$, since there is a unique rooted 1-manifold of diameter x for each $0 < x \leq \infty$. The space $\mathcal{M}^1_2$ is then naturally homeomorphic to the set

$$S = \{ (x, y) \mid 0 < x < \infty,\ 0 \leq y \leq x \text{ or } x = \infty,\ 0 \leq y < \infty \},$$

where x is the diameter of M and y is the distance between the base points $p, q \in M$. Both the left and right projections of $\mathcal{M}^1_2$ onto $\mathcal{M}^1$ are then the first coordinate projection $S \to (0, \infty]$.

Let μ be a σ-finite Borel measure on $\mathcal{M}^d$. From μ, we define two associated Borel measures μ_l and μ_r on $\mathcal{M}^d_2$, by setting

$$\mu_l(S) = \int_{(M,p)\in\mathcal{M}^d} \mathrm{vol}_M \{ q \in M \mid (M, p, q) \in S \} \, d\mu$$

$$\mu_r(S) = \int_{(M,q)\in\mathcal{M}^d} \mathrm{vol}_M \{ p \in M \mid (M, p, q) \in S \} \, d\mu,$$

whenever S is a Borel subset of $\mathcal{M}^d_2$. Sometimes, we abbreviate the above as

$$d\mu_l(M, p, q) = d\mathrm{vol}_M(q)\, d\mu(M, p), \quad d\mu_r(M, p, q) = d\mathrm{vol}_M(p)\, d\mu(M, q).$$

Definition 2.3. We say μ is unimodular if $\mu_l = \mu_r$. When μ is a probability measure, a μ-random element of $\mathcal{M}^d$ is a unimodular random manifold (URM).

Unimodular measures were first studied in the context of rooted graphs rather than rooted Riemannian manifolds (see [16], [57], [9]). In these works, the equality of μ_l and μ_r is phrased via the mass transport principle, which is the definition we gave in the introduction (Definition 18). Lewis Bowen [25] has previously considered unimodularity in the general context of metric-measure spaces; however here we restrict ourselves to the Riemannian setting.

When $d = 1$, any Borel probability measure on $\mathcal{M}^d = (0, 1]$ is unimodular, as the measures μ_l and μ_r are obtained by integrating μ against twice the Lebesgue measure on the fiber $[0, x] \subset S$ over $x \in (0, 1]$.

For an alternative definition, note that there is an involution

$$i : \mathcal{M}^d_2 \to \mathcal{M}^d_2, \quad i(M, p, q) = (M, q, p) \tag{6}$$

and that $i(\mu_l) = \mu_r$. It follows that a measure μ on $\mathcal{M}^d$ is unimodular if and only if either/both of μ_l and μ_r are i-invariant.

As mentioned in the introduction, any finite volume Riemannian d-manifold M determines a unimodular measure μ, obtained by pushing forward the Riemannian volume vol_M under the map

$$M \to \mathcal{M}^d, \quad p \mapsto (M, p).$$

In this case, the measures μ_l and μ_r are both obtained by pushing forward the product measure $\mathrm{vol}_M \times \mathrm{vol}_M$ on $M \times M$ to $\mathcal{M}^d_2$, so μ is unimodular. Also, $\mu \mapsto (\mu_l, \mu_r)$ is weak* continuous, so the space of unimodular measures on $\mathcal{M}^d$ is closed. So, more unimodular measures can be constructed as weak* limits.

Here are some other constructions of unimodular measures on $\mathcal{M}^d$.

Example 2.4 (Regular covers). Suppose that $\pi : N \to M$ is a regular Riemannian covering map and M has finite volume. Then there is a map

$$i : M \to \mathcal{M}^d, \quad i(p) = (N, q), \text{ where } \pi(q) = p.$$

Here, the point is that the isometry class of (N, q) depends only on the projection $p = \pi(q)$, since any two q with the same projection differ by a deck transformation. The push forward of vol_M under i is a probability measure μ on $\mathcal{M}^d$ that is supported on manifolds isometric to N. For an alternative construction of μ, choose a fundamental domain F for the projection $\pi : N \to M$ and push forward the measure $\mathrm{vol}_N / \mathrm{vol}_N(F)$ via the map $q \in F \mapsto (N, q)$.

To see that μ is unimodular, let $\mathcal{N}_2 = N \times N / \Gamma$, where Γ is the group of deck transformations of π, which acts diagonally on $N \times N$. We can identify $\mathcal{N}_2$ with $F \times N$, and give it the measure $\mathrm{vol}_N / \mathrm{vol}_N(F) \times \mathrm{vol}_N$. Then the map

$$\mathcal{N}_2 = N \times N / \Gamma \to N \times N / \mathrm{Isom}(N) \subset \mathcal{M}^d_2$$

is measure preserving, where we consider $\mathcal{M}^d_2$ with μ_l. On $\mathcal{N}_2$, the involution $(p, q) \mapsto (q, p)$ is measure preserving, since for each $\gamma \in \Gamma$ the composition

$$F \times \gamma F \subset F \times N = \mathcal{N}_2 \xrightarrow{(p,q) \mapsto (q,p)} \mathcal{N}_2 = F \times N$$

is just given by $(x, y) \mapsto (\gamma^{-1}(y), \gamma^{-1}(x))$. So, as this involution on $\mathcal{N}_2$ descends to (6) on $\mathcal{M}^d_2$, the measure μ_l is invariant under (6), so μ is unimodular.

Example 2.5 (Restriction to saturated subsets). A subset $B \quad \mathcal{M}^d$ is saturated if whenever $(M; p) \in B$ and $q \in M$, then $(M; q) \in B$ as well. Note that saturated Borel subsets of $\mathcal{M}^d$ form a -algebra, $\mathcal{S}$.

If is unimodular and B is a saturated Borel subset of $\mathcal{M}^d$, then j_B is unimodular as well, since $(_B)_l$ and $(_B)_r$ are just the restrictions of $_l = _r$ to the set of all $(M; p; q) \in \mathcal{M}_2$ with $(M; p) \in B$.

Finally, let X be a complete Riemannian manifold with a transitive isometry group $\mathrm{Isom}(X)$. Up to rooted isometry, the choice of root in X is irrelevant, so we will denote the corresponding point in $\mathcal{M}^d$ by X as well.

Proposition 2.6. If a Riemannian manifold X has transitive isometry group, the atomic probability measure $_X$ supported on $X \in \mathcal{M}^d$ is unimodular if and only if $\mathrm{Isom}(X)$ is a unimodular Lie group.

Proof. Fix a point $x_0 \in X$ and let $K < G = \mathrm{Isom}(X)$ be the stabilizer of x_0, so that we can identify $X = G/K$. Then $(_X)_l$ is supported on

$$X_2 := G \backslash X \quad X$$

where G acts diagonally. Since G acts transitively, the natural map

$$K \backslash X \xrightarrow{[x] \to [(x_0; x)]} X_2: \tag{7}$$

is a homeomorphism. With respect to this identication, $(_X)_l$ is just the push forward of the Riemannian measure on X to $K \backslash X$.

Since $K \backslash X = K \backslash G/K$, the identication (7) can also be written as

$$K \backslash G/K \xrightarrow{KgK \to [(x_0; g(x_0))]} X_2:$$

Conjugating, the involution $[(p; q)] \to [(q; p)]$ on X_2 becomes the inversion map

$$\bar{i} : K \backslash G/K \to K \backslash G/K; \quad \bar{i}([g]) = [g^{-1}]:$$

So, with (6) in mind, we want to show that the natural measure on $K \backslash G/K$ is $\bar{i}$-invariant if and only if G is unimodular.

Integrate against the (unique) right K-invariant probability measures on the bers of $G \to G/K$; this gives a left Haar measure on G. Then G is unimodular if and only if is invariant under the inversion map

$$i : G \to G; \quad i(g) = g^{-1}:$$

By denition, can be expressed as an integral

$$= \int \quad _{KgK} \, d; \; KgK \in K \backslash G/K$$

where the ber measure $_{KgK}$ is the unique probability measure on KgK that is K-biinvariant. The action of $i : G \to G$ permutes the $_{KgK}$, which implies that is i-invariant if and only if the factor measure is $\bar{i}$-invariant.

As semisimple groups are unimodular, their symmetric spaces X satisfy the assumptions above. One can also see directly that the atomic measure $_X$ is unimodular when X is a model space of constant curvature, i.e. when $X = \mathbb{R}^n; \mathbb{H}^n$ or S^n. The measures $(_X)_l$ and $(_X)_r$ are supported on the subset $X_2 \quad \mathcal{M}^d_2$ consisting of isometry classes of doubly pointed manifolds $(X; p; q)$. Here, these $(X; p; q)$ are classified up to isometry by $d(p;q)$, which is symmetric in $p;q$. So, the involution i in (6) is the identity, and therefore preserves $(_X)_l;(_X)_r$.

2.1. The No-Core Principle.

This trivial, yet useful, consequence of unimodularity was mentioned in §1.3.4.

Theorem 1.15 (The No-Core Principle). Suppose that is a unimodular probability measure on $\mathcal{M}^d$ and that $f : \mathcal{M}^d \to \{0;1\}$ is a Borel function. Then for -almost every element $(M;p) \in \mathcal{M}^d$, we have

$$0 < \mathrm{vol}_M\{q \in M \mid f(M;q) = 1\} < \infty \implies \mathrm{vol}(M) < \infty:$$

It is very important here that is a probability measure, and not just -finite. Otherwise, one could take a fixed Riemannian manifold N with infinite volume and no symmetries, and any finite (nonzero) volume subset $B \subset M$, and define $f(M;p) = 1$ if there is an isometry $M \to N$ that takes p into B. As N has no symmetries, the map $N \to \mathcal{M}^d;\ q \to (N;q)$ is an embedding, so vol_N pushes forward to a -finite unimodular measure $_N$ on $\mathcal{M}^d$, and the pair f and $_N$ violates the statement of the theorem.

Proof. If the theorem fails, then for some $C > 0$ the set of all $(M;p)$ such that

$$0 < \mathrm{vol}_M\{q \in M \mid f(M;q) = 1\} < C \text{ and } \mathrm{vol}(M) = \infty \tag{8}$$

has positive -measure. This set is saturated and Borel, so we may assume after restriction and rescaling that is supported on it, as in Example 2.5.

Well, by unimodularity we know that

$$\int_{(M;p)\in\mathcal{M}^d}\int_{q\in M} f(M;p)\, d\mathrm{vol}_M\, d = \int_{(M;p)\in\mathcal{M}^d}\int_{q\in M} f(M;q)\, d\mathrm{vol}_M\, d:$$

On the right, the integrand is at most C -almost surely, by (8). So, the right-hand side is finite. Therefore, the integrand on the left is finite -almost surely. So, for -a.e. $(M;p)$, we have $f(M;p) = 0$ by (8). This implies that the left side is zero. So, the integrand on the right side is zero -a.e., contradicting (8).

2.2. Unimodularity and IRSs.

In previous work with Bergeron, Gelander, Nikolov, Raimbault and Samet, the authors studied the following group theoretic analogue of URMs, see [2]. Let G be a locally compact, second countable topological group and let Sub_G be the space of closed subgroups of G, endowed with the Chabauty topology, see §A.4.

Definition 2.7. An invariant random subgroup (IRS) of G is a random element of Sub_G whose law is a Borel measure invariant under conjugation by G. (In an abuse of notation, we will often refer to the law itself as an IRS.)

Invariant random subgroups supported on discrete groups of unimodular G satisfy a useful group-theoretic unimodularity property. Fix a Haar measure for G. If H is a discrete subgroup of G, then pushes forward locally to a Radon measure $_H$ on the coset space $H\backslash G$. Let Cos_G be the set of cosets of closed subgroups of G, endowed with its Chabauty topology.

Theorem 2.8 (Biringer-Tamuz [21]). Assume G is unimodular, and μ is a Borel probability measure on Sub_G such that μ-a.e. $H \in \mathrm{Sub}_G$ is discrete. Then μ is an IRS if and only if for every Borel function $f : \mathrm{Cos}_G \to \mathbb{R}$, we have

$$\int_{\mathrm{Sub}_G}\int_{H\backslash G} f(Hg)\, d\lambda_H(Hg)\, d\mu(H) = \int_{\mathrm{Sub}_G}\int_{H\backslash G} f(g^{-1}H)\, d\lambda_H(Hg)\, d\mu(H).$$

As noted in [21], a more aesthetic version of the equality above is

$$\int_{H\in \mathrm{Sub}_G}\int_{g\in H\backslash G} f(Hg)\, d\lambda_H\, d\mu = \int_{H\in \mathrm{Sub}_G}\int_{g\in G/H} f(gH)\, d\lambda^H\, d\mu;$$

where λ^H is the measure on G/H obtained by locally pushing forward λ. In other words, the 'right' measure obtained on Cos_G by viewing it as the space of right cosets and then integrating the natural invariant measures on $H\backslash G$ against μ is the same as the analogous 'left' measure on Cos_G. The version given in the theorem is that which we will use here, though.

Suppose now that our unimodular G acts isometrically and transitively with compact stabilizers on a Riemannian d-manifold X, and write $X = G/K$, where K is a compact subgroup of G. A $(G;X)$-manifold is a quotient $H\backslash X$, where $H < G$ acts freely and properly discontinuously. Let

$$\mathrm{Sub}_G^{df} = \{H \in \mathrm{Sub}_G \mid H \text{ acts freely and properly discontinuously on } X\};$$

and let $\mathcal{M}^{(G;X)} \subset \mathcal{M}^d$ be the space of all pointed $(G;X)$-manifolds.

Proposition 2.9 (URM vs. IRS, transitive case). The continuous map

$$\Psi : \mathrm{Sub}^{df}_G \to \mathcal{M}^{(G;X)}, \quad \Psi(H) = (H\backslash X; [id])$$

induces a (weak* continuous) map

$$\Psi_* : \left\{ \text{IRSs } \mu \text{ of } G \text{ with } \mu(\mathrm{Sub}_G^{df}) = 1 \right\} \to \left\{ \text{URMs } \mu \text{ with } \mu(\mathcal{M}^{(G;X)}) = 1 \right\}.$$

If X is simply connected and $G = \mathrm{Isom}(X)$ is the full isometry group of X, then Ψ_* is a weak* homeomorphism.

Recall that a unimodular random manifold (URM) is a random element of a unimodular probability measure μ on $\mathcal{M}^d$. However, we routinely abuse terminology by calling measures IRSs, and we will similarly call μ itself a URM. Note that Ψ_* is surjective, but is not in general injective, since conjugating H by an element of K does not change its image. Continuity of Ψ_* follows from Proposition 3.10 and Lemma 3.7 of [3].

Proof. Fix a Haar measure λ on G normalized so that $\pi_*\lambda = \mathrm{vol}_X$, where $\pi : G \to G/K = X$ is the projection. If $H \in \mathrm{Sub}^{df}_G$ and $\pi_H : H\backslash G \to H\backslash X$ is the natural projection, it follows that $(\pi_H)_*\lambda_H = \mathrm{vol}_{H\backslash X}$.

Unimodularity of the image. Let μ be an IRS with $\mu(\mathrm{Sub}^{df}_G) = 1$. We must show that the measures $(\Psi_*\mu)_l$ and $(\Psi_*\mu)_r$ in Definition 2.3 are equal. So, let $f : \mathcal{M}^d_2 \to \mathbb{R}$ be a Borel function. We then define a new function

$$\tilde{f} : \mathrm{Cos}_G^{df} \to \mathbb{R}; \quad \tilde{f}(Hg) = f(H\backslash X; [id]; [g]);$$

where Cos_G^{df} is the set of cosets of subgroups $H \in \mathrm{Sub}_G{}^{df}$. We then compute:

$$\int_{\mathcal{M}^d_2} f\, d(\Psi_*\mu)_l = \int_{(H\backslash X;[id])\in\mathcal{M}^d}\int_{[g]\in H\backslash X} f(H\backslash X;[id];[g])\, d\mathrm{vol}_{H\backslash X}\, d(\Psi_*\mu)$$

= Z Z f̃(Hg) d_H d;
Z^{H 2 Sub_G} Z^{Hg 2 H n G}
= f̃((g ¹Hg)g ¹) d_H d;
Z^{H 2 Sub_G} Hg 2 H n G Z
= f(g ¹HgnX; [id]; [g ¹]) dvol_{HnX} d()
Z^{(HnX;[id]) 2 M^d} Z^{[g] 2 HnX}
= f(HnX; [g]; [id]) dvol_{HnX} d()
Z^{(HnX;[id]) 2 M^d} [g] 2 HnX
= f d()_r : M_2^d

Here, the rst equation is the denition of d()_l, keeping in mind that it is enough to integrate over rooted manifolds of the form (HnX; [id]). The second and forth equations follow from the normalization of the Haar measure, while the third is Theorem 2.8. The fth equation reects the fact that (HnX; [g]; [id]) and (g ¹HgnX; [id]; [g ¹]) are isometric as doubly rooted manifolds.

The case of the full isometry group. Assume now that X is simply connected and G = Isom(X) is the full isometry group of X.

We rst analyze the bers of : Sub^df_G ! M^(G;X). Conjugate subgroups of G give isometric X-quotients, and if two subgroups H; H⁰ 2 Sub_G^df are conjugate by an element of K, then the pointed manifolds (HnX; [id]) and (H⁰nX; [id]) are isometric. Conversely, as K is the full group of isometries of X xing [id], any based isometry of quotients lifts to a K-conjugacy of subgroups, so we have:

(9) bers of : Sub^df_G ! M^(G;X) ! K-conjugacy classes in Sub^df_G:

Injectivity. Let is an IRS with (Sub_G^df) = 1. By Rohlin's disintegration theorem (see [88, Theorem 6.2]), disintegrates as an integral

= Z_{(M;p) 2 M^(G;X)} _(M;p) d();

where _(M;p) is a Borel probability measure on the preimage ¹(M; p). The K-action by conjugation on Sub^df_G leaves the -bers invariant (9) and preserves , so it must preserve -a.e. ber measure _(M;p). As each -ber is a K-homogeneous space, each _(M;p) is just the push forward of the unique Haar probability measure on K. Therefore, can be recovered by integrating the canonical measures _(M;p) against . So, is injective.

Surjectivity. If is an URM with (M^(G;X)) = 1, dene a measure on Sub_G by

(10) = Z_{(M;p) 2 M^(G;X)} _(M;p) d;

where as above each _(M;p) is the unique K-invariant probability measure on ¹(M; p). Then () = , and we claim that is an IRS.

Our strategy will be to use the unimodularity of to establish the equality in Theorem 2.8 (the mass transport principle for IRSs). Consider the map

_2 : Cos_G^df ! M_2^(G;X); _2(Hg) = (HnX; [id]; [g]);

where $\mathcal{M}_2^{(G,X)}$ is the space of doubly rooted X-manifolds. The fibers of π_2 are exactly the $K \times K$-orbits $(Hg)^{K \times K}$ in Cos_G^{df}, where the action is defined as

$$(11) \qquad K \times K \curvearrowright \mathrm{Cos}_G, \qquad (k,k') \cdot Hg := (kHk^{-1})kgk'.$$

We define $\eta_{(M,p,q)}$ to be the unique $K \times K$-invariant probability measure on the fiber $\pi_2^{-1}(M,p,q)$, i.e. the push forward of the Haar measure on $K \times K$ under the conjugation action.

Claim 2.10. For $(M,p) \in \mathcal{M}^{(G,X)}$, we have $\int_H \nu_H \, d\eta_{(M,p)}(H) = \int \eta_{(M,p,q)} \, d\mathrm{vol}_M(q)$.

We will prove the claim below, but first we use it to prove that ν is an IRS, by deriving the equality in Theorem 2.8 from the unimodularity of μ. Suppose that $f : \mathrm{Cos}_G^{df} \longrightarrow \mathbb{R}$ is a Borel function, and define a new function

$$(\pi_2)_* f : \mathcal{M}_2^{(G,X)} \longrightarrow \mathbb{R}, \quad (\pi_2)_* f(M,p,q) := \int_{Hg \in \pi_2^{-1}(M,p,q)} f(Hg) \, d\eta_{(M,p,q)}.$$

We first compute the left side of the equality in Theorem 2.8.

$$\begin{aligned}
&\int_{\mathrm{Sub}_G} \int_{H \backslash G} f(Hg) \, d\nu_H(Hg) \, d\nu(H) \\
&= \int_{(M,p) \in \mathcal{M}^d} \int_{H \in \pi^{-1}(M,p)} \int_{H \backslash G} f(Hg) \, d\nu_H(Hg) \, d\eta_{(M,p)} \, d\mu \\
&= \int_{(M,p) \in \mathcal{M}^d} \int_{q \in M} \int_{Hg \in \pi_2^{-1}(M,p,q)} f(Hg) \, d\eta_{(M,p,q)} \, d\mathrm{vol}_M \, d\mu \\
&= \int_{\mathcal{M}_2^{(G,X)}} (\pi_2)_* f \, d\mu_l
\end{aligned}$$

and using a similar argument, we compute the right side:

$$\begin{aligned}
&\int_{\mathrm{Sub}_G} \int_{H \backslash G} f(g^{-1}H) \, d\nu_H(Hg) \, d\nu(H) \\
&= \int_{(M,p) \in \mathcal{M}^d} \int_{q \in M} \int_{Hg \in \pi_2^{-1}(M,p,q)} f(g^{-1}H) \, d\eta_{(M,p,q)} \, d\mathrm{vol}_M \, d\mu \\
&= \int_{(M,p) \in \mathcal{M}^d} \int_{q \in M} \int_{Hg \in \pi_2^{-1}(M,q,p)} f(Hg) \, d\eta_{(M,q,p)} \, d\mathrm{vol}_M \, d\mu \\
&= \int_{\mathcal{M}_2^{(G,X)}} (\pi_2)_* f \, d\mu_r.
\end{aligned}$$

So, the unimodularity of μ implies that ν is an IRS.

Weak* homeomorphism. Finally, recall that (10) defines an inverse for Φ. Weak* continuity of the inverse will follow if we show that the map

$$\mathcal{M}^{(G,X)} \longrightarrow \mathcal{P}(\mathrm{Sub}_G), \quad (M,p) \longmapsto \eta_{(M,p)}$$

is continuous, where $\mathcal{P}(\mathrm{Sub}_G)$ is the space of Borel probability measures on Sub_G, considered with the weak* topology. However, $\eta_{(M,p)}$ is the unique K-invariant measure on $\pi^{-1}(M,p)$, and if $(M_i,p_i) \to (M,p)$, we can pass to a subsequence so that $\eta_{(M_i,p_i)}$ converges. The limit must be supported on $\pi^{-1}(M,p)$, and is K-invariant since its approximates are, so must be $\eta_{(M,p)}$.

Finally, we promised to prove Claim 2.10 during the proof above:

Proof of Claim 2.10. Let $(M,p,q) \in \mathcal{M}_2^{(G,X)}$ and let $Hg \in \mathrm{Cos}_G$ such that $\pi_2(Hg) = (M,p,q)$. By (11), we have a commutative diagram

$$\begin{array}{ccc} K\times K & \xrightarrow{\quad (k,k')\,\mapsto\,(kHk^{-1})kgk' \quad} & \mathrm{Cos}_G \\ \Big\downarrow \pi_l & & \Big\downarrow \pi_r \\ K\times K \,/\, N_K(H)\times K & \xrightarrow{\quad [(k,k')]\,\mapsto\, kHk^{-1} \quad} & \mathrm{Sub}_G, \end{array}$$

where $\pi_r : \mathrm{Cos}_G \to \mathrm{Sub}_G$, $\pi(Hg) = H$. As we had previously defined $\mu_{(M,p)}$ as the K-invariant probability measure on $H^K \subset \mathrm{Sub}_G$, the diagram shows that

$$(\pi_r)_*(\nu_{(M,p,q)}) = \mu_{(M,p)}.$$

The Haar probability measure on $K\times K$ disintegrates under π_l as an integral of invariant probability measures on the cosets of $N_K(H)\times K$ against the pushforward measure on $K\times K / N_K(H)\times K$. Here, the coset $(k,1)N_K(H)\times K$ has a measure invariant under its stabilizer, which is $N_K(kHk^{-1})\times K$. This disintegration pushes forward to a π_r-disintegration of $\nu_{(M,p,q)}$:

$$\nu_{(M,p,q)} = \int_{F\in \mathrm{Sub}_G} \nu^F_{(M,p,q)} \, d\mu_{(M,p)}, \tag{12}$$

where $\nu^F_{(M,p,q)} = 0$ unless F is a conjugate of H, in which case $\nu^F_{(M,p,q)}$ is an invariant probability measure on the $N_K(F)\times K$-orbit in Cos_G obtained by intersecting $\pi_2^{-1}(M,p,q)$ with $F\backslash G$.

Now fix $(M,p) \in \mathcal{M}^{(G,X)}$, let $H \in \pi^{-1}(M,p)$ and fix an isometric identification of $(H\backslash G/K,[id])$ with (M,p). The fibers of the composition

$$\begin{array}{ccccc} H\backslash G & \longrightarrow & M & \longrightarrow & \mathcal{M}_2^{(G,X)} \\ Hg & \longmapsto & [g] \in H\backslash G/K & \longmapsto & (H\backslash G/K,[id],[g]). \end{array}$$

are exactly the $N_K(H)\times K$-orbits in $H\backslash G$. As ν_H is invariant under the action of $N_K(H)\times K$, it disintegrates as an integral of invariant probability measures on these orbits against its pushforward under the composition. Under the first map, ν_H pushes forward to vol_M, so we may write instead:

$$\nu_H = \int_{q\in M} \nu^H_{(M,p,q)} \, d\mathrm{vol}_M. \tag{13}$$

Combining Equations (12) and (13), we can now prove the claim:

$$\begin{aligned} \int_{\mathrm{Sub}_G} \nu_H \, d\mu_{(M,p)} &= \int_{\mathrm{Sub}_G}\int_M \nu^H_{(M,p,q)} \, d\mathrm{vol}_M(q)\, d\mu_{(M,p)} = \\ &= \int_M \int_{\mathrm{Sub}_G} \nu^H_{(M,p,q)} \, d\mu_{(M,p)} \, d\mathrm{vol}_M(q) \\ &= \int_M \nu_{(M,p,q)} \, d\mathrm{vol}_M(q). \end{aligned}$$ $\square$

We now construct URMs from IRSs of discrete groups. Suppose that G is a discrete group that acts freely and properly discontinuously on a Riemannian d-manifold X and that the quotient $G\backslash X$ has finite volume. There is a map

$$\text{an IRS } \mu \text{ of } G \;\longmapsto\; \bar{\mu} \text{ a probability measure on } \mathcal{M}^{(G,X)};$$

where a $\bar\mu$-random element of $\mathcal{M}^d$ has the form $(H\backslash X, [x])$, where we first take $x \in X$ to be an arbitrary lift of a random point in $G\backslash X$ and then choose $H \in \mathrm{Sub}_G$ μ-randomly. The conjugation invariance of μ makes the measure $\bar\mu$ well-defined despite the arbitrary choice of lift. Alternatively, consider

$$B = (\mathrm{Sub}_G \times X)/\Gamma, \quad \text{where } (H, x) \mapsto (H^{\gamma^{-1}}, \gamma x).$$

Then B is a Sub_G-bundle over $G\backslash X$, and each of its fibers has an identification with Sub_G that is canonical up to conjugation. So, as μ is conjugation invariant there is a well-defined probability measure μ_B on B obtained as the integral of μ on each fiber against the (normalized) Riemannian volume of $G\backslash X$. The map

$$\mathrm{Sub}_G \times X \longrightarrow \mathcal{M}^d, \quad (H, x) \longmapsto (H\backslash X, [x])$$

factors through the Γ-action to a map $B \longrightarrow \mathcal{M}^d$, and $\bar\mu$ is the push forward of μ_B under this map.

Proposition 2.11 (IRS $\Rightarrow$ URM, discrete case). If μ is an IRS of Sub_G, then $\bar\mu$ is an URM.

Proof. Pick a Borel fundamental domain $D \subset X$, i.e. a Borel subset such that

1) $\mathrm{vol}_X(D \cap gD) = 0$ for every $g \in G$,
2) $\mathrm{vol}_X(X \setminus \cup_{g\in G} gD) = 0$.

It follows that $\mathrm{vol}_X = \sum_{g\in G} g(\mathrm{vol}_D)$ and moreover that if $H \in \mathrm{Sub}_G$, then

$$\mathrm{vol}_{H\backslash X} = \sum_{Hg \in H\backslash G} (\pi \circ g)(\mathrm{vol}_D); \tag{14}$$

where $\pi : X \longrightarrow H\backslash X$ is the quotient map. We let $\hat\mu$ be the push forward to $\mathcal{M}^d$ of $\mu \times \mathrm{vol}_D$ under the function

$$\mathrm{Sub}_G \times D \longrightarrow \mathcal{M}^d, \qquad (H, x) \longmapsto (H\backslash X, [x]).$$

This $\hat\mu$ is the scale by $\mathrm{vol}_X(D)$ of our $\bar\mu$ above. For simplicity of notation, we show that $\hat\mu$ is unimodular instead.

We must show that $\hat\mu_l = \hat\mu_r$, so let $f : \mathcal{M}^d_2 \longrightarrow \mathbb{R}$ be a Borel function. We lift f to a function $\tilde f : \mathrm{Cos}_G \longrightarrow \mathbb{R}$ by letting

$$\tilde f(Hg) = \int_{(x,y)\in D} f\,(H\backslash X, [x], [gy])\; d\mathrm{vol}_D^2;$$

Note that $[gy] \in H\backslash X$ only depends on the coset Hg. We now compute:

$$\begin{aligned}
&\int_{\mathcal{M}^d_2} f\, d\hat\mu_l \\
&= \int_{(H\backslash X,[x])\in \mathcal{M}^d} \int_{[y]\in H\backslash X} f\,(H\backslash X, [x], [y])\; d\mathrm{vol}_{X/H}\, d\hat\mu \\
&= \int_{(H\backslash X,[x])\in \mathcal{M}^d} \sum_{Hg\in H\backslash G} \int_{x\in D} f\,(H\backslash X, [x], [gy])\; d\mathrm{vol}_D\, d_H\, d\hat\mu \qquad (15)\\
&= \int_{H\in \mathrm{Sub}_G} \sum_{Hg\in H\backslash G} \int_{(x,y)\in D^2} f\,(H\backslash X, [x], [gy])\; d\mathrm{vol}_D^2\, d_H\, d\mu \\
&= \int_{H\in \mathrm{Sub}_G} \sum_{Hg\in H\backslash G} \tilde f(Hg)\, d_H\, d\mu
\end{aligned}$$

$$(16)\qquad = \int_{H\in \mathrm{Sub}_G} \sum_{Hg\in H\backslash G} \tilde f\big((g^{-1}Hg)g^{-1}\big)\, d\lambda_H\, d\mu$$

$$= \int_{H\in \mathrm{Sub}_G} \sum_{Hg\in H\backslash G} \int_{(x,y)\in D^2} f\big(g^{-1}Hg\backslash X, [x], [g^{-1}y]\big)\, d\mathrm{vol}^2_D\, d\lambda_H\, d\mu$$

$$(17)\qquad = \int_{H\in \mathrm{Sub}_G} \sum_{Hg\in H\backslash G} \int_{(x,y)\in D^2} f\big(H\backslash X, [gx], [y]\big)\, d\mathrm{vol}^2_D\, d\lambda_H\, d\mu$$

$$(18)\qquad = \int_{(H\backslash X,[y])\in \mathcal{M}^d} \int_{[x]\in H\backslash X} f\big(H\backslash X, [x], [y]\big)\, d\mathrm{vol}_{X/H}\, d\hat\mu$$

$$= \int_{\mathcal{M}^d_2} f\, d\hat\mu_r.$$

Above, (15) and (18) follow from (14), while (16) is Proposition 2.8. Line (17) uses the fact that $\big(g^{-1}Hg\backslash X, [x], [g^{-1}y]\big)$ and $(H\backslash X, [gx], [y])$ are isometric as doubly rooted manifolds.

3. Measures on Riemannian foliated spaces

A foliated space with tangential dimension d is a separable metrizable space X that has an atlas of charts of the form

$$\phi : U \to L \times Z,$$

where each $L \subset \mathbb{R}^d$ is open and each Z is a separable, metrizable space. Tran-sition maps must preserve and be smooth in the horizontal direction, with partial derivatives that are continuous in the transverse direction. The horizontal fibers piece together to form the leaves of X. See [34] and [79] for details. A foliated space X is Riemannian if each of its leaves has a smooth, complete Riemannian metric, and if these metrics vary smoothly in the transverse direction, in the sense that the charts ϕ can be chosen so that if $t_n \to t \in Z$, the induced Riemannian metrics g_{t_n} on L converge smoothly to g_t.

We are interested in measures on a Riemannian foliated space X that are formed by integrating vol against a 'transverse measure'. To this end, suppose that $\mathcal{U} = \{(U_i, \phi_i)\}$ is a countable atlas of charts as above and let $Z = \cup Z_i$ be the as-sociated 'transverse space'. An invariant transverse measure on X is a σ-finite measure on Z that is invariant under the holonomy groupoid of $\mathcal{U}$. Here, the holo-nomy groupoid is that generated by homeomorphisms between an open subset of some Z_i and an open subset of some Z_j that are defined by following the leaves of the foliation (see [34]). The reader can verify that if $\mathcal{U}$ and $\mathcal{U}'$ are countable atlases associated to a foliated space X, there is a 1-1 correspondence between the invariant transverse measures of $\mathcal{U}$ and those of $\mathcal{U}'$.

If λ is an invariant transverse measure on a Riemannian foliated space X, one can locally integrate the Riemannian measure vol against λ to give a measure μ on X, specified by writing $d\mu = \mathrm{vol}\, d\lambda$. For a precise definition, let $\phi : U \to L \times Z$ be a foliated chart and define a measure μ on U by the formula

$$\mu(E) = \int_{x\in Z} \mathrm{vol}\big(E \cap \phi^{-1}(L \times x)\big)\, d\lambda.$$

Then if ffg is a partition of unity subordinate to our atlas, we dene

$$= \sum f \; :$$

Using holonomy invariance, one can check that the measure does not depend on the chosen partition of unity.

Measures on a Riemannian foliated space X satisfying d = vol d for some are usually called completely invariant. Actually, complete invariance just ensures that when is disintegrated locally along the leaves of the foliation, Lebesgue measure is recovered in the tangent direction; that is, a transverse measure is automatically holonomy invariant whenever the measure vol d on the ambient space is well-dened (see [40]).

Theorem 3.1. Suppose that X is a Riemannian foliated space and is a -nite Borel measure on X. Then the following are equivalent:

1) is completely invariant.
2) is unimodular, as dened in Equation (4) of §1.3.
3) lifts uniformly to a measure ~ on the unit tangent bundle T^1X that is invariant under geodesic ow, see (19) below.

If the leaves of X have bounded geometry[6], then 1) { 3) are equivalent to

4) $\int_X \mathrm{div}(Y)\, d = 0$ for every vector eld Y on X with integrable leaf-wise divergence.
5) $\int_X f g\, d = \int_X f g\, d$ for all continuous functions $f; g : X \to \mathbb{R}$ that are C^2 on each leaf of X.

The meaning of 3) was explained in the introduction, but briey, the leaf-wise unit tangent bundle T^1X maps onto X, and the bers T^1_pX are round spheres. If $!_p$ is the Riemannian measure on the ber T_xX, then we can dene

$$d\tilde{} = \; !_p \; d; \qquad (19)$$

so $d\tilde{}$ is a measure on T^1X. The geodesic ows on the unit tangent bundles of the leaves of X then piece together to a well-dened geodesic ow on T^1X, and 3) says that this ow leaves the measure ~ invariant.

As discussed in the introduction, this result may be particularly interesting to those familiar with unimodularity in graph theory. Condition 3) is similar to the 'involution invariance' characterization of unimodularity of Proposition 2.2 in [9]. Also, in analogy with 5), the 'graphings' of [50] can be characterized via the self adjointness of their Laplacian. See [60] for one direction; the other direction follows from the arguments in [73, Proposition 18.49].

The equivalence 1) , 4) is well-known as a consequence of work of Lucy Garnett [51], and a version with slightly dierent hypotheses on the foliated space appears in a recent paper of Catuogno{Ledesma{Runo [36]. However, we include the very brief proof below.

Proof of 1) =) 2). Suppose that d = vol d for some . Let $_i$:

$$U_i \to L_i \times Z_i; \quad i = 1; 2;$$

---

[6] This condition is needed only in 5) =) 1), in order to invoke a theorem of Garnett [51]. It means that there is some uniform K such that every point x 2 X lies in a smooth coordinate patch for its leaf that has derivatives up to order 3 bounded by K, see [51].

be two foliated charts for X and assume that there is a homeomorphism : Z_1 ! Z_2 in the holonomy pseudogroup. We rst check that $_r$ = $_l$ on the set X $_{;\,;}$ X X dened by

$$X_{1;2;} = f(x_1; x_2) \; 2 \; U_1 \; U_2 \, j \, _i(x_i) = (l_i; z_i) \text{ where } (z_1) = z_2 g: \text{For a}$$

subset S $X_{1;2;}$, we then calculate

$$(_l S) = Z \; Z \; Z \quad 1_S(x; y) \text{ dvol dvol d}$$

$$Z^{z_1 2 Z_1} \; Z^{x 2_1^{\;1}(L_1 f z_1 g)} \; Z^{y 2_2^{\;1}(L_2 (z_1))}$$

$$= \quad 1_S(x; y) \text{ dvol dvol d} \quad ()$$

$$z_2 2 Z_2 \quad y 2_2^{\;1}(L_2 f z_2 g) \quad x 2_1^{\;1}(L_1 f \;{}^1(z_2) g)$$

$$= {}_r(S);$$

Above, () follows from a change of variables, the invariance of under the holo-nomy pseudogroup and Fubini's Theorem. Now, both of the measures $_r$ and $_l$ are supported on the equivalence relation R X X of the foliation. However, we claim that R can be covered by a countable number of the Borel subsets X $_{;\,;}$, which will prove the claim. First, the separability of X guarantees that X X can be covered by a countable number of open sets U_1 U_2 with $_i$: U_i ! L_i Z_i foliated charts. If a pair of points with coordinates $(l_i; z_i)$ 2 L_i Z_i; i = 1; 2 determines an element of $(U_1 U_2) \setminus R$; then $z_2 = (z_1)$ for some holonomy map . The set of germs of holonomy maps taking a given z_1 2 Z_1 into Z_2 is countable, so as Z_1 is separable, a countable number of domains and ranges of holonomy maps suce to cover $(U_1 \; U_2) \setminus R$.

Proof of 2) =) 3). Suppose is a unimodular measure on X and let ~ be the induced measure on the foliated space T^1X. Each leaf of T^1X is the unit tangent bundle of a leaf of X and the tangential Riemannian metric is the Sasaki metric. The Riemannian volume on each leaf of T^1X is then the berwise product of vol with the Lebesgue measures $!_x$ on the tangent spheres T^1X_x.

First, note that ~ is unimodular. For T^1X T^1X bers over X X and $d\tilde{}_l = d!_y \, d!_x \, d_l$ while $d\tilde{}_r = d!_y \, d!_x \, d_r$; so the fact that $_l = {}_r$ implies that $\tilde{}_l = \tilde{}_r$. Geodesic ow $_t$ lifts to a map $_t$ = $(_t; id)$ on T^1X T^1X; as Liouville measure is geodesic ow invariant, the measure $\tilde{}_r$ is clearly $_t$-invariant. As ~ is unimodular, this implies that $\tilde{}_l$ is $_t$-invariant. But under the rst coordinate projection T^1X T^1X ! T^1X, $\tilde{}_l$ pushes forward to ~ and $_t$ descends to $_t$, so it follows that ~ is $_t$ invariant.

Before proving that 3) =) 1), we need the following lemma.

Lemma 3.2. Suppose that U is an open subset of a Riemannian manifold M and denote the geodesic ow on T^1M by g_t. Let be a Borel measure on U and let d~ = $!_p$ d be the lifted measure on T^1U, as in (19). Suppose that for all Borel subsets S T^1U and all t 2 R with $g_t(S)$ T^1U we have $\tilde{}(g_t(S)) = \tilde{}(S)$. Then is a scale of the Riemannian measure on U.

The rst part of this proof was shown to us by Nir Avni.

Proof. We rst prove that is absolutely continuous with respect to the Riemann-ian measure on U, so let S U be a set of -measure zero, and let $\tilde{S}$ be the set of all pairs (p; v) 2 T^1U where p 2 S. After subdividing S into countably many pieces, we may assume that there is some > 0 such that for all p 2 S, we have B(p;) U and < $\mathrm{inj}_M(p)$, where $\mathrm{inj}_M(p)$ is the injectivity radius of M at p.

Choose a probability measure supported in (0;) that is absolutely continuous with respect to Lebesgue measure. For each p 2 S, dene a map

$$\tilde{}_p : T^1_p U \quad (0;) \quad ! \quad T^1U; \quad {}_p(v; \tilde{t}) = g_t(v);$$

Note that as B(p;) U the image of the map does in fact lie in U. We then dene a measure ${}_p$ on U via the formula

$${}_p = (\ {}_p)(!_p \ \tilde{} \ );$$

where : T^1U ! U is the projection map. As < $\mathrm{inj}_M(p)$, the map $\tilde{}_p$ is a dieomorphism onto its image, so the pushforward ${}_p$ is absolutely continuous with respect to the Riemannian measure on U. Then we have

$$(20) \qquad (S) = \tilde{}(\tilde{S}) = \int_{t2(0;)} \tilde{}(g_t(\tilde{S}))\, d$$

$$= \int_{t2(0;)} \int_{p2U} \int_{v2T^1_pU} 1_{g_t(\tilde{S})}\, d!_p\, dd$$

$$= \int_{p2U} \left(\int_{t2(0;)} \int_{v2T^1_pU} 1_{g_t(\tilde{S})}\, d!_p\, d \right) d$$

$$= \int_{p2U} !_p \quad f(v; t)\ j\ g_t(v) \ 2 \ Sg \ \tilde{} \quad d$$

$$(21) \qquad = \int_{p2U} {}_p(S)\ d$$

$$(22) \qquad = \int_{p2U} 0\ d \ = \ 0:$$

Here, (20) comes from the g_t-invariance of ~ and the fact that is a probability measure. Equation (21) follows since $\tilde{S}$ consists of all unit tangent vectors lying above points of S and the projection is injective on the image of $\tilde{}_p$. Finally, equation (22) is just the fact that ${}_p$ is absolutely continuous to the Riemannian measure on U, with respect to which S has measure 0. This shows that is absolutely continuous with respect to , which also implies that ~ is absolutely continuous with respect to the Liouville measure ~.

To show that is a scalar multiple of ; consider the commutative triangle

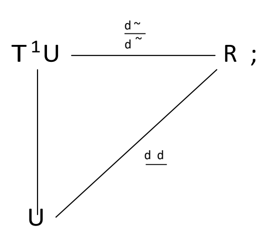

where $\frac{d\tilde{}}{d\tilde{}}$ and $\frac{d}{d}$ are the Radon-Nikodym derivatives. Since any two points $p; q$ in U can be joined by a geodesic in M, there are unit tangent vectors $(p; v)$ and $(q; w)$ with $g_t(p; v) = (q; w)$ for some t. But since both ~ and ~ are geodesic ow invariant,

$$\frac{d}{d}(p) = \frac{d\tilde{}}{d\tilde{}}(p; v) = \frac{d\tilde{}}{d\tilde{}}(q; w) = \frac{d}{d}(q):$$

It follows that is a scalar multiple of .

Proof of 3) =) 1). Let : U ! L Z be a foliated chart for X. The restriction j_U then disintegrates as $dj_U = {}_z\, d;$ where

is the pushforward of under the projection L Z ! Z, and
each ${}_z$ is a Borel probability measure on L fzg.

The map z ! ${}_z$ is Borel, in the sense that for any Borel B L Z we have that z ! ${}_z(B)$ is Borel.

Consider now the foliated chart ~ : T^1U ! T^1L Z for T^1X. The lifted measure $d\tilde{} = !_p\, d$ then disintegrates as $!_{(z;l)} d_z\, d$. As ~ is invariant under the geodesic ow $g_t : T^1X$! T^1X, it follows that for -almost all z 2 Z, the measure $!_{(z;l)} d_z$ is invariant under the geodesic ow of Lfzg, regarded as an open subset of its leaf in X. Thus, by Lemma 3.2, the probability measure ${}_z = \frac{\mathrm{vol}_z}{\mathrm{vol}(Lfzg)}$ for -almost all z 2 Z. Since this is true within every U, there is a holonomy invariant transverse measure ^{0}, dened locally by $^0 = \mathrm{vol}_z(L\ fzg)$, with $= \mathrm{vol}\, d^0$: This proves the claim.

Proof of 1 =) 4. Assume that d = vol d and that Y is a continuous vector-eld on X with integrable divergence on each leaf. Decomposing Y using a partition of unity, we may assume that Y is supported within some compact subset of the domain of a foliated chart : U ! L Z. Then

$$\int_X \mathrm{div}(Y)\, d = \int_{z 2 Z} \int_{Lfzg} \mathrm{div}(Y)\, d\mathrm{vol}_z\, d$$
$$= \int_{z 2 Z} 0\, d = 0;$$

by the divergence theorem applied to each leaf L fzg.

Proof of 4) =) 5). We compute:

$$\int_X f\ g\, d = \int_X f\ \mathrm{div}(rg)\, d$$
$$= \int_X \mathrm{div}(frg)\quad hrf; rgi\, d$$
$$= \int_X hrf; rgi\, d;$$

by condition 4). As this is symmetric in f and g, condition 5) follows.

Proof of 5) $\Rightarrow$ 1). It follows immediately from 5) that $\int_X \Delta f \, d\mu = 0$ for every continuous $f : X \to \mathbb{R}$ that is C^2 on each leaf of X. In the terminology of Garnett [51], μ is harmonic. Using the bounded geometry condition, Garnett proves that in every foliated chart $\phi : U \to L \times Z$, a harmonic measure μ disintegrates as $d\mu = h(l, z) \, d\mathrm{vol}_z \, d\nu(z)$, where h is a positive leaf-wise harmonic function and ν is a measure on the transverse space Z. We must show that $h(\cdot, z)$ is constant for ν-almost every z.

If $f, g \in C(X)$ are continuous functions supported in some compact subset of U that are C^2 on each plaque $L \times \{z\}$, we have by 5) that

$$\int_{z \in Z} \int_{L \times \{z\}} f \, \Delta g \, h \, d\mathrm{vol}_z \, d\nu = \int_{z \in Z} \int_{L \times \{z\}} \Delta f \, g \, h \, d\mathrm{vol}_z \, d\nu$$

$$= \int_{z \in Z} \int_{L \times \{z\}} f \, \Delta (g \, h) \, d\mathrm{vol}_z \, d\nu$$

As f is arbitrary, this implies that $\Delta g \, h = \Delta (g \, h)$ on ν-almost every plaque $L \times \{z\}$. As g is arbitrary, $h(\cdot, z)$ must be constant for ν-a.e. z. $\square$

4. The foliated structure of $\mathcal{M}^d$

Let $\mathcal{M}^d$ be the space of isometry classes of pointed Riemannian manifolds (M, p), equipped with the smooth topology. The space $\mathcal{M}^d$ is separable and completely metrizable — we refer the reader to the appendix §A.1 for a detailed introduction to the smooth topology and a proof of this result.

4.1. Regularity of the leaf map.

When X is a d-dimensional Riemannian foliated space, there is a 'leaf map'

$$\mathcal{L} : X \to \mathcal{M}^d, \quad \mathcal{L}(x) = (L_x, x),$$

defined by mapping each point x to the isometry class of the pointed manifold (L_x, x), where L_x is the leaf of X containing x. We claim:

Proposition 4.1 (The leaf map is Borel). *If $U \subset \mathcal{M}^d$ is open, then $\mathcal{L}^{-1}(U) = \bigcup_{i \in \mathbb{N}} O_i \cap C_i$, where each O_i is open and each C_i is closed in X.*

In [68, Lemma 2.8], Lessa showed that the leaf map is measurable when the Borel σ-algebra of X is completed with respect to any Borel probability measure on X. The proof is a general argument that any construction in a Lebesgue space that does not use the axiom of choice is measurable, and uses the existence of an inaccessible cardinal. He remarks that a more direct investigation of the regularity of $\mathcal{L}$ can probably be performed, which is what we do here. We should also mention that Alvarez López and Candel [71] study the leaf map from a foliated space into the Gromov-Hausdorff space of pointed metric spaces, and have observed, for instance, that it is continuous on the union of leaves without holonomy. See also [10], where together with Barral Lijo, they study the leaf map into $\mathcal{M}^d$.

The key to proving Proposition 4.1 is the following slight extension of a result of Lessa [68, Theorems 4.1 & 4.3], which we prove in the appendix.

Theorem A.19. *Suppose X is a d-dimensional Riemannian foliated space in which $x_i \to x$ is a convergent sequence of points. Then $\mathcal{L}(x_i)$ is pre-compact in $\mathcal{M}^d$, and every accumulation point is a pointed Riemannian cover of $\mathcal{L}(x)$.*

There is a partial order on M^d, where $(N; q)$ $(M; p)$ whenever $(N; q)$ is a pointed Riemannian cover of $(M; p)$. With respect to , Theorem A.19 asserts an 'upper semi-continuity' of the leaf map. The degree of regularity of L indicated in Proposition 4.1 is exactly that of upper semicontinuous maps of between ordered spaces, so to get the same conclusion in our setting we must show a compatibility between and the smooth topology on M^d:

Lemma 4.2. Every point $(M; p)$ 2 M^d has a basis of neighborhoods U such that the following properties hold for each U 2 U:

1) there is no $(N; q)$ $(M; p)$ such that $(N; q)$ 2 @U,
2) if $(N^0; q^0)$ $(N; q)$ $(M; p)$ and $(N^0; q^0)$ 2 U, then $(N; q)$ 2 U.

Proof. In §A.2, we dene the open k^{th}-order $(R;)$-neighborhood of $(M; p)$,

$$N^k_{R;1;}(M; p);$$

to be the set of all $(N; q)$ such that there is an embedding $f : B_M(p; R)$! N with $f(p) = q$ such that $D^k f : T^k U$! $T^k N$ is locally ^{0}-bilipschitz with respect to the iterated Sasaki metrics on the 1-neighborhood of the zero section in $T^k U$, where $1 <\ ^0 <$. Any sequence of these neighborhoods is a basis around $(M; p)$ as long as ! 1 and $R; k$! 1, and we will show that when $; R; k$ are chosen appropriately then these neighborhoods satisfy the conditions of the lemma.

The subset C M^d of pointed covers of $(M; p)$ is compact: if $R > 0$ is given the uniform geometry bounds on $B(p; R)$ M lift to any cover, see Denition A.3, so Theorem A.4 gives pre-compactness of C M^d, and C is closed in M^d since Arzela-Ascoli allows one to take a limit of covering maps. Now x some $R > 0$. If $(N_i; q_i)$ 2 C is a convergent sequence, the isometry type of $B(q_i; R)$ N_i is eventually constant. So by compactness, the R-ball around the base point takes on only nitely many isometry types within C.

Arzela-Ascoli's theorem implies that when forming the closure of $N^k_{R;}(M; p)$, we just allow $^0 =$. So, the boundaries $@N^k_{R;}(M; p)$ are disjoint for distinct values of . As there only nitely many isometry types of, say, 2R-balls around the base point in pointed covers of $(M; p)$, there can be only nitely many < 2 such that there is a cover of $(M; p)$ in $@N^k_{R;}(M; p)$. So, the rst condition in the lemma is satised as long as we choose < 2 to avoid these points.

To illustrate which neighborhoods $N^k_{R;}(M; p)$ satises the second condition of the lemma, we need the following:

Claim 4.3. Fix $(M; p)$ 2 M^d. Then for all R in an open, full measure subset of $R_{>0}$, there is some > 1 such that whenever

$$^0 : (N^0; q^0) \;\; ! \;\; (M; p)$$

is a pointed Riemannian covering and

$$f : B(p; R) \;\; ! \;\; N^0$$

is a locally -bilipschitz embedding with $f(p) = q^0$, then ^{0} is injective on $f(B(p; R))$.

Proof. If not, there is a sequence indexed by i such that $_i$! 1, but $_i$

$$: (N_i; q_i^0) \;\; ^C! \;\; (M; p)$$

is not injective on $f_i(B(p; R))$. In the limit, we obtain a Riemannian cover

$$^0 : (N^0; q^0) \;\; !(M; p);$$

such that there is a pointed isometry

$$f : B(p, R) \to B(q', R) \subset N',$$

but where π' is not injective on $\overline{B(q', R)}$. Here, the non-injectivity persists in the limit since the distance between points of (N'_i, q'_i) with the same projection to (M, p) is bounded below by the injectivity radius of (M, p), which is positive on the compact subset of (M, p) in which we're interested.

The map $\pi' \circ f : B(p, R) \to M$ is an isometry fixing p, so it extends to an embedding $F : \overline{B(p, R)} \to M$. As long as R is a (generalized) regular value for the (nonsmooth) function $d(p, \cdot)$ on M, a full measure open condition [84], the inclusion $B(p, R) \hookrightarrow \overline{B(p, R)}$ is a homotopy equivalence, see [38, Isotopy Lemma 1.4]. So in this case, the map F takes $\pi_1(\overline{B(p, R)}) = \pi_1(B(p, R))$ into the π'-image of $\pi_1(N', q')$. Hence F lifts to an isometry $\overline{B(p, R)} \to B(q', R) \subset N'$, by the lifting criterion. Since F is an embedding, this contradicts that π' is non-injective on $\overline{B(q', R)}$.

As long as ϵ, R are chosen according to Claim 4.3, $N^k_{R,\epsilon}(M, p)$ satisfies the second condition of the lemma. For if

$$(N', q') \to (N, q) \to (M, p), \qquad (N', q') \in N^k_{R,\epsilon}(M, p),$$

then there is a map $f : B(p, R) \to N'$ as above, so $\pi' : (N', q') \to (M, p)$ is injective on $f(B(p, R))$. In particular, the covering map

$$\pi : (N', q') \to (N, q)$$

is also injective there, so the composition

$$\pi \circ f : B(p, R) \to N'$$

is an embedding. As $\pi \circ f$ inherits the same Sasaki-bilipschitz bounds that f has, this shows that $(N, q) \in N^k_{R,\epsilon}(M, p)$ as well.

Therefore, for any k, almost every $R > 0$, and ϵ sufficiently close to 1, the neighborhood $N^k_{R,\epsilon}(M, p)$ satisfies both conditions of the lemma. As these neighborhoods form a basis for the topology of M^d at (M, p), we are done.

Using the lemma, we now complete the proof of Proposition 4.1. Recall that X is a d-dimensional Riemannian foliated space and

$$L : X \to M^d, \quad x \mapsto (L_x, x)$$

is the leaf map. We want to show that for each open $U \subset M^d$, the preimage $L^{-1}(U) = \cup_{i \in \mathbb{N}} O_i \cap C_i$, where each O_i is open and each C_i is closed in X.

It suffices to check this when U is chosen as in Lemma 4.2. If $L^{-1}(U)$ does not have the form $\cup_{i \in \mathbb{N}} O_i \cap C_i$, there is a point $x \in L^{-1}(U)$ and a sequence

$$x_i \in \overline{L^{-1}(U)} \setminus L^{-1}(U), \quad x_i \to x \in L^{-1}(U).$$

Passing to a subsequence, we may assume by Theorem A.19 that

$$L(x_i) \to (N, q) = L(x).$$

Note that as $L(x_i) \notin U$ for all i, we have $(N, q) \notin U$ as well.

Each x_i is the limit of some sequence $(y_{i,j})$ in $L^{-1}(U)$, and Theorem A.19 implies that after passing to a subsequence, we have that for each i,

$$L(y_{i,j}) \to (Z_i, z_i) = L(x_i) \text{ as } j \to \infty.$$

Now xing $R > 0$, since the manifolds $L(x_i)$ converge in M^d, the R-balls around their base points have uniformly bounded geometry, as in Denition A.3. These geometry bounds lift to pointed covers, so are inherited by the $(Z_i; z_i)$. So by Theorem A.4, after passing to a subsequence we may assume that

$$(Z_i; z_i) \to (N^0; q^0) \in M^d:$$

Moreover, as $(Z_i; z_i)$ $L(x_i)$ for each i, we have

$$(N^0; q^0) = \lim(Z_i; z_i) \ \lim L(x_i) = (N; q);$$

simply by taking a limit of the covering maps. Remembering now that $(Z_i; z_i)$ was dened as the limit of $L(y_{i;j})$ as $j \to 1$, if we choose for each i some large $j = j(i)$ and abbreviate $y_i = y_{i;j(i)}$, then

$$L(y_i) \to (N^0; q^0)$$

as well. However, by construction we have $L(y_i) \in U$, so $(N^0; q^0) \in \overline{U}$.

The rst part of Lemma 4.2 implies that $(N^0; q^0) \in U$, and then the second part shows $(N; q) \in U$. This is a contradiction, as we said above that $(N; q) \notin U$.

4.2. Resolving singularities in M^d. n this section, we will assume that d 2. We saw in Example 2.2 that $M^1 = (0; 1]$ is completely understood; the reader is encouraged to think through the proofs of our results when $d = 1$ on his/her own.

M^d is not a naturally foliated space: although the images of the maps

$$M \to M^d; \quad p \to (M; p)$$

partition M^d as would the leaves of a foliation, these maps are not always injective and their images may not be manifolds. However, the following theorem, discussed in §1, shows that there is a way to desingularize M^d so that the theory of unimodular measures becomes that of completely invariant measures.

Theorem 1.8 (Desingularizing M^d). If is a completely invariant probability measure on a Riemannian foliated space X, then pushes forward under the leaf map to a unimodular probability measure on M^d.

Conversely, there is a Polish Riemannian foliated space P^d such that any -nite unimodular measure on M^d is the push forward under the leaf map of some completely invariant measure on P^d. Moreover, for any xed manifold M, the preimage of $f(M; p) \mid p \in Mg$ M^d under the leaf map is a union of leaves of P^d, each of which is isometric to M.

As a corollary of this and Theorem 3.1, we have the following theorem, which we also discussed in the introduction.

Theorem 1.9. A -nite Borel measure on M^d is unimodular if and only if the lifted measure ~ on T^1M^d is geodesic ow invariant.

Proof. By Theorem 1.8, is the push forward of a completely invariant measure on a Riemannian foliated space $: X \to M^d$. (Taking $X = P^d$.) By Theorem 3.1, the induced measure ~ on T^1X is invariant under geodesic ow. Now, the leafwise derivative $D : T^1X \to T^1M^d$ is geodesic ow equivariant, so the push forward measure D~ = ~ is geodesic ow invariant.

The rst assertion of Theorem 1.8 is easy to prove. If

$$R = f(x; y) \in X \times X \mid x; y \text{ lie on the same leaf}g;$$

then the measures $\mu_{l;r}$ on $X \times X$ are supported on R and push forward to $\nu_{l;r}$ under the natural map $R \to M^d$. By Theorem 3.1, $\nu_l = \nu_r$, so $\mu_l = \mu_r$.

The idea for the 'conversely' statement is to use Poisson processes to obstruct the symmetries of these manifolds, converting M^d into a foliated space P^d. To do this, we will recall some background on Poisson processes, define P^d and show how to translate between measures on M^d and on P^d, and then verify that P^d is a Riemannian foliated space.

If M is a Riemannian d-manifold, the Poisson process of M is the unique probability measure $\mathbb{P}_M$ on the space of locally finite subsets $D \subset M$ such that

1) if $A_1, \ldots, A_n$ are disjoint Borel subsets of M, the random variables that record the sizes of the intersections $D \cap A_i$ are independent,
2) if $A \subset M$ is Borel, the size of $D \cap A$ is a random variable having a Poisson distribution with expectation $\mathrm{vol}_M(A)$.

For a finite volume subset $A \subset M$ and $n \in \mathbb{N}$, we have (cf. [41, Example 7.1(a)])

$$\text{(23)} \qquad \mathrm{Prob}\begin{pmatrix}\text{for } (x_1,\ldots,x_n)\in A^n, \text{ we have } D\cap A=\{x_1,\ldots,x_n\},\\ \text{given that } D\cap A \text{ has } n \text{ elements.}\end{pmatrix} = d\,\mathrm{vol}_M^n(x_1,\ldots,x_n).$$

In other words, if D is chosen randomly, the elements of $D \cap A$ are distributed within A independently according to vol_M.

We refer the reader to [41] for more information on Poisson processes. In this text, they are not introduced on Riemannian manifolds, but for measures on $\mathbb{R}^d$ that are absolutely continuous with respect to Lebesgue measure. However, as the Poisson process behaves naturally under restriction and disjoint union, it is 'local', and can be defined naturally for manifolds. In fact, the Poisson process really only depends on the Riemannian measure on M, and not on the topology of M. Since M is isomorphic as a measure space to the (possibly infinite) interval $(0, \mathrm{vol}(M))$, see [85], one really only needs to understand the usual Poisson process on $\mathbb{R}_+$, as that of an interval is just its restriction.

When M is a Riemannian d-manifold, let $\mathcal{F}M$ be its orthonormal frame bundle, the bundle in which the fiber over $p \in M$ is the set of orthonormal bases for T_pM. If we regard $\mathcal{F}M$ with the Sasaki-Mok metric [78], then

1) when $f : M \to M$ is an isometry, so is its derivative $Df : \mathcal{F}M \to \mathcal{F}M$,
2) the Riemannian measure $\mathrm{vol}_{\mathcal{F}M}$ is obtained by integrating the Haar probability measure on each fiber $\mathcal{F}_pM \cong O(d)$ against vol_M.

Lemma 4.4. *If M is a Riemannian d-manifold, $d \geq 2$, then $\mathrm{Isom}(M)$ acts essentially freely, with respect to the Poisson measure $\mathbb{P}_{\mathcal{F}M}$, on the set of nonempty locally finite subsets of $\mathcal{F}M$.*

The subset $\emptyset \subset \mathcal{F}M$ is fixed by $\mathrm{Isom}(M)$ and has $\mathbb{P}_{\mathcal{F}M}$-probability[7] $e^{-\mathrm{vol}\,\mathcal{F}M}$. So if M has finite volume, we must exclude $\emptyset$ in the statement of the lemma.

Also, if $M = S^1$, then after choosing an orientation, every subset $\{e_1, e_2\} \subset \mathcal{F}M$, where e_1, e_2 have opposite orientations, is stabilized by an involution of M. Since S^1 is compact, two-element subsets of $\mathcal{F}S^1$ appear with positive $\mathbb{P}_{\mathcal{F}S^1}$-probability, so the statement of the lemma fails for 1-manifolds.

Proof. The Lie group $\mathrm{Isom}(M)$ acts freely on $\mathcal{F}M$, so any nonempty subset $D \subset \mathcal{F}M$ that is stabilized by a nontrivial element $g \in \mathrm{Isom}(M)$ has at least two points.

[7] Via the measure isomorphism $\mathcal{F}M \to (0, \mathrm{vol}(\mathcal{F}M))$, this is just the probability that there are no points in the interval $(0, \mathrm{vol}(\mathcal{F}M))$, under the usual Poisson process on $\mathbb{R}_+$.

As the action is proper, its orbits are properly embedded submanifolds, so unless one is a union of components of FM, all orbits have vol_{FM}-measure zero. In this case, the vol_{FM}-probability of selecting two points from the same orbit is zero, so μ_{FM}-a.e. $D \subset FM$ has trivial stabilizer, by Equation (23).

So, we may assume from now there is an orbit of $\mathrm{Isom}(M) \curvearrowright FM$ that is a union of components of FM. (The frame bundle FM has either 1 or 2 components, depending on whether M is orientable.) Then $\mathrm{Isom}(M)$ acts transitively on 2-planes in TM, so M has constant sectional curvature. As $\mathrm{Isom}(M)$ also acts transitively on TM, M is either $S^d, \mathbb{RP}^d, \mathbb{R}^d$ or $\mathbb{H}^d$.

If $D \subset FM$ is stabilized by some nontrivial $g \in \mathrm{Isom}(M)$, it has at least two points e_1, e_2, and we can consider the images $g(e_1), g(e_2)$. Either e_1, e_2 are exchanged by g, or one is sent to the other, which is sent to something new, or both elements are sent to new elements of D. As the elements of a random D are distributed according to vol_{FM}, by (23), it suffices to prove that for $(e_1, \dots, e_4) \in FM^4$, the following are vol_{FM}-measure zero conditions:

1) $g(e_1) = e_2$, $g(e_2) = e_1$, for some $g \in \mathrm{Isom}(M)$,
2) $d(e_1, e_2) = d(e_2, e_3)$,
3) $d(e_1, e_2) = d(e_3, e_4)$.

An isometry that exchanges two frames must be an involution, since its square fixes a frame. So, for 1) we want to show that the probability of selecting frames $e_1, e_2 \in FM$ that are exchanged by an involution is zero. The point is that in each of the cases $M = S^d, \mathbb{RP}^d, \mathbb{R}^d$ or $\mathbb{H}^d$, an involution exchanging $p, q \in M$ leaves invariant some geodesic $\gamma : [0,1] \to M$ joining p, q, and then exchanges $\gamma'(0) \in T_pM$ with $\gamma'(1) \in T_qM$. So, after fixing a frame $e_1 \in FM_p$, the frames in FM_q that are images of e_1 under involutions form a subset of FM_q of dimension at most that of $O(d-1)$, which has zero Haar measure inside of $FM_q = O(d)$. Integrating over q, we have that for a fixed e_1, the probability that a frame $e_2 \in FM$ is the image of e_1 under an involution is zero. Integrating over e_1 finishes the proof of part 1).

For 2), note that for a fixed $e_2 \in FM$, the function

$$d(e_2, \cdot) : FM \to \mathbb{R}$$

pushes forward vol_{FM} to a measure on $\mathbb{R}$ that is absolutely continuous with respect to Lebesgue measure – its RN-derivative at $x \in \mathbb{R}$ is the $(\dim(FM)-1)$-dimensional volume of the metric sphere around e_2 of radius x. So, if e_1 and e_3 are chosen against vol_{FM}, the distances $d(e_1, e_2)$ and $d(e_2, e_3)$ will be distributed according to a measure absolutely continuous to Lebesgue measure on $\mathbb{R}^2$, so will almost never agree. The proof that 3) is a measure zero condition is similar.

We now show how to convert $\mathcal{M}^d$ into a foliated space by introducing Poisson processes on the frame bundles of each Riemannian d-manifold. Let

$$\mathcal{P}^d_{\mathrm{all}} = \left\{ (M, p, D) \;\middle|\; \begin{smallmatrix} M \text{ a complete Riemannian } d\text{-manifold,} \\ p \in M, \text{ and } D \subset FM \text{ a closed subset} \end{smallmatrix} \right\} \Big/ \sim,$$

where $(M, p, D) \sim (M', p', D')$ if there is an isometry $\phi : M \to M'$ with $\phi(p) = p'$ whose derivative $d\phi$ takes D to D'. There is a Polish smooth-Chabauty topology on $\mathcal{P}^d_{\mathrm{all}}$ obtained from the smooth topology on $\mathcal{M}^d$ and the Chabauty topology on the subsets D, see §A.5. Now consider the subset

$$\mathcal{P}^d = \{ (M, p, D) \in \mathcal{P}^d_{\mathrm{all}} \mid \nexists \text{ an isometry } \phi : M \to M \text{ with } d\phi(D) = D \}.$$

The subset P^d is G, since P^d_{all} n P^d = $[_{n2N} F_n$; where F_n is the set of all (M; p; D) such that there is an isometry : M ! M with

$$d(D) = D \text{ and } 1{=}n \ \ d_M(p; (p)) \ \ n;$$

here, F_n is closed by the Arzela Ascoli theorem. Hence, by Alexandrov's theorem, P^d is a Polish space. Note that P^d is dense in P^d_{all}, since any Riemannian manifold can be perturbed to have no nontrivial isometries.

Theorem 4.5. P^d has the structure of a Polish Riemannian foliated space, where (M; p; D) and (M^0; p^0; D^0) lie in the same leaf when there is an isometry

$$: M \ \ ! \ \ M^0; \ \ d(D) = D^0:$$

Assuming Theorem 4.5 for a moment, let's indicate how to transform a unimodular measure on M^d into a completely invariant measure on P^d, which will nish the proof of Theorem 1.8. Each ber of the projection

$$: P^d_{all} \ \ ! \ \ M^d; \ \ (M; p; D) \ \ ! \ \ (M; p)$$

is identied with a set of of closed subsets of F M , and this identication is unique up to isometry. The Poisson process on F M (with respect to the natural volume, e.g. that induced by the Sasaki-Mok metric) induces a measure $_{(M;p)}$ on ^{1}(M; p), supported on the Borel set of locally nite subsets. We call this measure the framed Poisson process on that ber, and by Lemma 4.4 we have

$$_{(M;p)}(P^d) = 1 \quad e^{\ \mathrm{vol} F M} > 0$$

for each (M; p) 2 M^d. Moreover, the map

$$M^d \ \ ! \ \ M(P^d); \ \ (M; p) \ \ ! \ \ _{(M;p)}$$

is continuous, where M(P^d) is the space of Borel measures on P^d, considered with the weak topology. (This follows from weak* continuity of the Poisson process associated to a space with a measure as the measure varies in the weak* topology, which is a consequence of (23).)

So, given a measure on M^d, we dene a measure ^ on P^d by

$$\text{(24)} \qquad \hat{} = \int_{(M;p)2M^d} \frac{_{(M;p)}}{1 \quad e^{\ \mathrm{vol} F M}} \, d:$$

The push forward of ^ under the projection to M^d is clearly , so we must only check that ^ is completely invariant.

Suppose that R P^d P^d is the leaf equivalence relation of P^d, i.e. the set of all pairs ((M; p; D); (M^0; p^0; D^0)) such that there is an isometry : M ! M^0 with d(D) = D^0. Each such pair determines a tuple (M^0; (p); p^0; D^0), a doubly rooted manifold together with a closed subset of its frame bundle, that is unique up to isometry. So, there is a map

$$R \ \ ! \ \ M^d_2; \ \ ((M; p; D); (M^0; p^0; D^0)) \ \ ! \ \ (M^0; (p); p^0);$$

where the ber over (M; p; q) 2 M^d_2 is canonical identied (up to isometry of M) with the set of closed subsets of F M on which Isom(M) acts freely. From the construction of ^, the measures $\hat{}_l$ and $\hat{}_r$ on R are obtained by integrating (rescaled) Poisson processes on each ber against $_l$ and $_r$. So, if $_l$ = $_r$, then $\hat{}_l$ = $\hat{}_r$, implying ^ is completely invariant by Theorem 3.1.

4.3. The proof of Theorem 4.5. The goal is to cover $\mathcal{P}^d$ by open sets U, together with homeomorphisms

$$h : \mathbb{R}^d \times Z \longrightarrow U,$$

where $Z = Z(U)$ is separable and metrizable, with transition maps

$$t : \mathbb{R}^d \times Z \longrightarrow \mathbb{R}^d \times Z', \quad t(x, z) = (t_1(x, z), t_2(z)),$$

where $t_1(x, z)$ is smooth in x, and where t_1, t_2 and all the partial derivatives $\frac{\partial t_1(x,z)}{\partial x_i}$ are continuous on $\mathbb{R}^d \times Z$. We also want the leaves of the foliation to be obtained by fixing M and $D \subset M$, and letting the base point $p \in M$ vary.

The construction of the charts will require some work — in outline, the idea is as follows. Given $X \in \mathcal{P}^d$, we show that on any small neighborhood $U \ni X$, there is an equivalence relation $\sim$ whose equivalence classes are obtained by taking some (M, p, D) and making slight variations of the base point p. Eventually, these equivalence classes will be the plaques $h(\mathbb{R}^d \times \{z\})$, and the quotient space $U/\sim$ will be the transverse space Z. That is, we will have a homeomorphism h to complete the following commutative diagram:

$$\begin{array}{ccc} \mathbb{R}^d \times Z & \xrightarrow{h} & U \\ \big\downarrow & & \big\downarrow \\ Z & \longrightarrow & U/\sim \end{array}$$

Of course, this cannot be done without a careful choice of U. It is not hard to choose U so that each $\sim$-equivalence class is a small disc of base points in some manifold M with a distinguished closed subset D. Shrinking U, we show that one can construct a continuously varying family of base frames, one for each of these M. Then, we use the Riemannian exponential maps associated to these frames to parameterize the $\sim$-classes, which allows us to identify them with $\mathbb{R}^d$ in a way that is transversely continuous.

Most of the proof involves constructing the base frames. Essentially, this is just a framed version of the following, which we can discuss now without introducing more notation. Given our neighborhood U of $X \in \mathcal{P}^d$, there is a section $s : U/\sim \longrightarrow U$ for the projection map with $s([X]) = X$. Briefly, the idea for this is as follows. Fix a metric d on $\mathcal{P}^d$, and for each $\sim$-class E, set

$$E_0 = \{p \in E \mid d(p, X) \leq 2d(E, X)\}.$$

The point $s(E) \in E$ is then defined to be the 'circumcenter' of E_0, when we regard E_0 as a small subset inside of a Riemannian M as above. (In what follows, this argument will be done for base frames, and the notation will be different.)

Before starting the proof in earnest, we record the following two lemmas, which should convince the reader that such a foliated structure is likely.

Lemma 4.6 (Leaf inclusions)**.** *Suppose that $X = (M, p, D) \in \mathcal{P}^d$, and define*

$$L : M \longrightarrow \mathcal{P}^d, \quad L(q) = (M, q, D).$$

Then L is a continuous injection, so the restriction of L to any precompact subset of M is a homeomorphism onto its image.

Proof. For continuity of L, note that if $q, q' \in \mathbb{R}^d$ and $|q - q'| = \epsilon$, there is a diffeomorphism f of $\mathbb{R}^d$ taking q to q', such that

1) f is supported in a 2-ball around q,
2) the k^{th} partial derivatives of f are all bounded by some C(k;), where C(k;) ! 0 as ! 0.

For instance, one can just x any bump function taking the origin to (1; 0; : : : ; 0) that is supported in a 2-ball around the origin, and conjugate it using appropriate Euclidean similarities. So if $q; q^0$ 2 M are close, we can choose an R^d-chart around them and transfer such an f to M to give an almost isometric map verifying that (M; q; D) and $(M; q^0; D)$ are close in P^d.

Injectivity of L comes from the denition of P^d, which requires that there are no nontrivial isometries of the pair (M; D). The statement about compact subsets of M follows from point set topology, as P^d is Hausdor.

We optimistically call the image of L the leaf through X 2 P^d, and write

$$L_X := L(M) \ P^d:$$

Via L, we will from now on regard L_X as a smooth Riemannian d-manifold that is topologically embedded (non-properly) in the space P^d. (The distance function on a leaf L_X will usually be written as d_{L_X}.) Under this identication, the point p becomes X and the subset D F M becomes a distinguished subset of the frame bundle of L_X. Note that the natural (manifold) topology on L_X is not the subspace topology induced from the inclusion L_X ,! P^d.

Lemma 4.7 (Chabauty convergence of leaves). Suppose X_i ! X in P^d. Then

1) for every point Y 2 L_X, there is a sequence Y_i 2 L_{X_i} with Y_i ! Y in P^d and $d_{L_{X_i}}(X_i; Y_i)$! $d_{L_X}(X; Y)$.
2) if Y_i 2 L_{X_i} with $d_{L_{X_i}}(X_i; Y_i) < 1$, then after passing to a subsequence we have Y_i ! Y 2 L_X, and $d_{L_{X_i}}(X_i; Y_i)$! $d_{L_X}(X; Y)$.

Note that there may be sequences Y_i 2 L_{X_i} that converge in P^d, but where $d_{L_{X_i}}(X_i; Y_i)$! 1, so the statements about distance above have content.

Proof. By denition of the convergence X_i ! X, see §A.5, after conjugating by the leaf inclusions, there are almost isometric maps

$$f_i : L_X \dashrightarrow L_{X_i}; \tag{25}$$

with $f_i(X) = X_i$ and where the derivatives Df_i pull back the distinguished subsets of the frame bundles L_{X_i} to a sequence that Chabauty converges to the distinguished subset of L_X. (Here, $\dashrightarrow$ means that f_i is dened on some ball around the base point, where the domains exhaust L_X as i ! 1, see A.5.)

For 1), the sequence $Y_i = f_i(Y)$ is dened for large i and converges to Y in P^d: one can use the f_i in the denition of convergence. For 2), the sequence $f_i^{1}(Y_i)$ is dened for large i, and is pre-compact in L_X by the condition on distances. So, after passing to a subsequence, it converges to some point Y 2 L_X, and we will have Y_i ! Y in P^d as well, again using the f_i in the denition of convergence.

4.3.1. Constructing the equivalence relation. For any > 0, dene a relation on P^d by letting

$$X \ Y \quad \text{if} \quad Y \ 2 \ L_X \text{ and } d_{L_X}(X; Y) \ :$$

Note that is reexive and symmetric, but we only have

$$X \ Y \text{ and } Y \ Z =) \ X \ _{+} \ Z;$$

rather than transitivity for a particular . In particular, the equivalence class of X with respect to the transitive closure of is exactly L_X.

However, each is transitive on suciently small subsets of P^d.

Lemma 4.8. If O 2 P^d, then for xed ; > 0, there is a neighborhood U 3 O on which the relations and agree. Hence, if U is chosen so that =2 on U, then is an equivalence relation on U.

Proof. Assuming this is not the case for some < , let X_i and Y_i be sequences in P^d that converge to O, with X_i Y_i; but X_i Y_i: Since X_i Y_i, the distance d_L $(X; Y)$ < 1. So, $_x$Lemma 4.7 2) applies, and we must have $d_{L_{X_i}}(X_i; Y_i)$! $d_{L_O}(O; O)$ = 0, violating that $d_{L_{X_i}}(X_i; Y_i)$ for all i.

From now on, we assume that all our neighborhoods U are small enough so that $_1$=$_2$, in which case $_1$ is an equivalence relation.

Lemma 4.9. The quotient topology on U= $_1$ is separable and metrizable.

Proof. Separability of U= $_1$ is immediate, as U is separable. Metrizability of U= $_1$ can be proved in the same way that we prove it for P^d in Section A.5. The only dierence is that when comparing two triples $(M; p; D)$ and $(M^0; p; D^0)$, we now let our maps f : $B_M(p; R)$! M^0 take p to any point within an -neighborhood of p^0. Ordinarily, such exibility would make it hard to establish a triangle inequality, but if U is suciently small, then such a map that realizes the distance between $(M; p; D)$ and $(M^0; p; D^0)$ will have small distortion, and using a limiting argument as in Lemma 4.8, one can show that in fact $f(p)$ must be (arbitrarily) close to p^0, so that the compositon of two such maps still takes base points within of base points.

4.3.2. A section of base frames. We describe here how to construct, for each equivalence class in U= $_1$, a base frame for the corresponding M. To make a precise statement, we need a framed version of our space P^d. Dene

$$FP^d = (M; e; D) \;\; \begin{array}{c} M \text{ a complete Riemannian d-manifold,} \\ e\,2\,FM; \text{ and } D\,F\,M \text{ a closed subset} \\ \text{such that 9 an isometry } f:_M\;!\,M;\; Df(D) = D \end{array} \;\; =\text{isometry:}$$

There is a natural Polish topology on FP^d, coming from the framed smooth topology on the pairs $(M; e)$ and the Chabauty topology on the subsets D; compare with Section A.5. We denote the natural projection by

$$: FP^d \;\; ! \;\; P^d:$$

Lemmas 4.6 and 4.7 have framed analogues. If $X = (M; e; D)$ 2 FP^d; then

$$L : FM \;\; ! \;\; FP^d; \;\; L(e^0) = (M; e^0; D)$$

is a continuous injection, and via L we will view its image L_X FP^d as (the frame bundle of) a smooth Riemannian manifold. Moreover, if

$$X_i = (M_i; e_i; D_i) \; ! \;\; X = (M; e; D)$$

in FP^d, then by denition, see §A.5, there are almost isometric maps

$$f_i : M \dashrightarrow M_i$$

such that the derivatives Df_i map e to e_i, and pull back (D_i) to a sequence of subsets of FM that Chabauty converges to D. Identifying frame bundles with the corresponding leaves in FP^d as above, these derivatives become maps

$$F_i : L_X \dashrightarrow L_{X_i} \tag{26}$$

These maps are almost isometric, in the sense that if L_X and L_{X_i} are equipped with the Sasaki-Mok metrics g_i induced by the Riemannian metrics on M_i and M, see [78], then we have $F_i^*(g_i) \to g$ in the C^1-topology. Using the F_i, one can prove a framed analogue of Lemma 4.7:

Lemma 4.10 (Chabauty convergence of leaves). If $X_i \to X$ in FP^d, then

1) for every point $Y \in L_X$, there is a sequence $Y_i \in L_{X_i}$ with $Y_i \to Y$ in FP^d and $d_{L_{X_i}}(X_i; Y_i) \to d_{L_X}(X; Y)$.
2) if $Y_i \in L_{X_i}$ with $d_{L_{X_i}}(X_i; Y_i) < \infty$, then after passing to a subsequence we have $Y_i \to Y \in L_X$, and $d_{L_{X_i}}(X_i; Y_i) \to d_{L_X}(X; Y)$.

Finally, the relation $\sim_1$ on P^d pulls back under to a relation on FP^d, which we will abusively call $\sim_1$ as well. On small subsets $V \subset FP^d$, $\sim_1$ is an equivalence relation on V, just as in the previous section.

Lemma 4.11. Every $F \in FP^d$ has a neighborhood V on which there is a continuous map $s : V \to FP^d$ with $s(F) = F$ that is constant on $\sim_1$-equivalence classes and satisfies $s(X) \sim_1 X$ for all $X \in V$.

So, s gives a continuous section for the map $V \to V/\sim_1$ near F.

The proof of Lemma 4.11 will occupy the rest of the section. The idea is as follows. We first select a distinguished subset of each equivalence class, essentially consisting of those points whose distances to F are at most twice the minimum distance to F in that equivalence class. We then show that these subsets vary continuously with the equivalence class. The desired section is constructed by always choosing the 'circumcenter' of the distinguished subset.

Fix a metric d_{FP^d} on FP^d and suppose V is a metric ball around F. Define

$$: V \to \mathbb{R}; \quad (X) = \inf\{d_{FP^d}(X^0; F) \mid X^0 \in V;\ X^0 \sim_1 X\}.$$

Each $\sim_1$-equivalence class in V is pre-compact in FP^d, since it is the image of a pre-compact subset under a continuous map $FM \to FP^d$ from the frame bundle of some manifold M. And if $X \in V$, the metric ball V contains all points of FP^d that are closer to F than X. So, the infimum is always achieved.

Claim 4.12. The map is continuous.

Proof. Suppose that in V we have

$$X_i = (M_i; e_i; D_i) \to X = (M; e; D)$$

and that $X_i^0 \sim_1 X_i$ realize the infimums defining (X_i). We can assume that $\sup d_{L_{X_i}}(X_i; X_i^0) \leq 2$, so after passing to a subsequence, Lemma 4.10 2) implies that X_i^0 converges in FP^d to a point $X^0 \in L_X$. However,

$$d_{FP^d}(F; X^0) = \lim_i d_{FP^d}(F; X_i^0) \leq \lim_i d_{FP^d}(F; X) = d_{FP^d}(F; X);$$

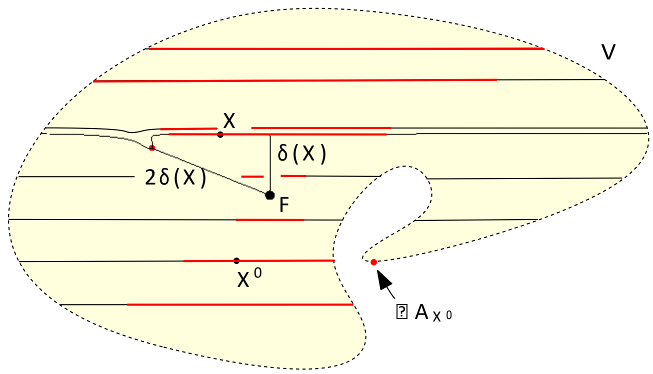

Figure 3. Each $A_X \subset [X] \cap V$ is drawn in red.

so as V is a metric ball around F that contains X, we have $X' \in V$. So, X' can be used in the definition of $\delta(X)$, implying that

$$\lim_i \delta(X_i) = \lim_i d_{FP^d}(F; X_i')^{C} = d_{FP^d}(F; X')\ \ \delta(X):$$

The reverse inequality is proved similarly: we assume that $X' \simeq_1 X$ realizes the infimum defining $\delta(X)$, then we reroot the X_i using Lemma 4.10 1) to produce elements $X_i' \in V$ with $X_i \simeq_1 X_i'$ and $X_i' \to X'$, and use the continuity of the distance function d_{FP^d}.

Moving toward the definition of the map s, we first define a map that selects a small subset of each $\simeq_1$-equivalence class in V. If $X \simeq_1 F$, set

$$A_X = \{F\}:$$

Otherwise, if $X \not\simeq_1 F$ define

$$A_X = \overline{\{Y \in V \mid Y \simeq_1 X \text{ and } d_{FP^d}(Y;F) < 2\delta(X)\}} \subset FP^d:$$

Note that by continuity of d_{FP^d} and positivity of $\delta(X)$, the set A_X is always nonempty. (However, if $X \simeq_1 F$ then $\delta(X) = 0$, so this latter definition of A_X would give the empty set, which is the reason we set $A_X = \{F\}$ in that case.) Moreover, A_X is compact: for if $X = (M; e; D)$ and

$$L : FM \to FP^d; \quad L(e') = (M; e'; D)$$

is the framed analogue of the leaf inclusion map of Lemma 4.6, then the conditions $L(e') \simeq_1 X$ and $d_{FP^d}(L(e'); F) < 2\delta(X)$ define a pre-compact subset of M, whose closure is the (compact) preimage $L^{-1}(A_X)$.

Claim 4.13 (Chabauty continuity of $X \to A_X$). After possibly shrinking V, we have that if $X_i \to X$ in V, then $A_{X_i} \to A_X$ in the Chabauty topology: every accumulation point in FP^d of a sequence $Z_i \in A_{X_i}$ lies in A_X, and every point of A_X is the limit of a sequence $Z_i \in A_{X_i}$.

Figure 3 indicates two issues related to this continuity.

1) The shape of V may be a problem: in the figure, the indicated point of $A_{X'}$ cannot be approached from 'below' within the red subsets. This is resolved by shrinking V.

2) It is important to dene A_X as the closure of a set dened by a strict inequality, as opposed to a set dened by a non-strict inequality. For the leftmost point in the gure that is equivalent to X has distance exactly $2(X)$ from F, but cannot be approached from above by red points.

Proof. Pick V^0 small enough so that for each X 2 V^0, the closed $2(X)$-ball around X 2 FP^d is contained in the interior of V. This is possible since (X) ! 0 as X ! F. Note that with V^0 so chosen, the condition that Y 2 V, rather than just in FP^d, is superuous in the denition of A_X.

Assume that X_i ! X in V^0. The fact that every accumulation point in FP^d of a sequence Z_i 2 A_{X_i} lies in A_X follows immediately from Lemma 4.10 2) and the continuity of d_{FP^d} and .

If Y 2 A_X, pick a sequence Y_i $_1$ X with $d_{FP^d}(Y_i; F) < 2(X)$ such that Y_i ! Y, as in the denition of A_X. Passing to subsequences of (X_i) and (Y_i) and using Lemma 4.10 2), pick for each i some Z_i $_1$ X_i with

$$d_{FP^d}(Y_i; Z_i) < \frac{1}{i}: \tag{27}$$

Note that we can assume Z_i 2 V^0, simply because the latter is open. As d_{FP^d} and are continuous on FP^d, for large i we have

$$d_{FP^d}(Z_i; X_i) < (X_i);$$

so Z_i 2 A_{X_i}. But Z_i ! Y, so we are done.

The goal now is dene a map $s : V$! FP^d by taking $s(X)$ to be the 'circumcenter' of A_X, in the following sense.

Lemma 4.14 (Circumcenters). Let M be a Riemannian manifold, let p 2 M and $R = \mathrm{inj}_M(p)$. Suppose that the sectional curvature of M is bounded above by on $B(p; R)$, and let $R^0 = \min\{R; \frac{}{16}\}$. If A $B(p; R^0)$, the function

$$q \; 2 \; M \; ! \; \sup\{d_M(q; a) \mid a \; 2 \; A\} \tag{28}$$

has a unique minimizer (A) 2 M, called the circumcenter of A.

Here, $\mathrm{inj}_M(p)$ is the injectivity radius of M at p, and the notation for (A) reects that it is the center of a minimal radius ball containing A.

Proof. As A $B(p; D)$ is pre-compact, the supremum above is always realized on the closure of A. Also, pre-compactness implies that the function (28) is continuous and proper, so it must have at least one minimum.

Suppose $q_1 = q_2$ both realize the minimum, which we'll say is D, and let q be the midpoint of the unique minimal geodesic connecting $q_1; q_2$. Since D is the minimum, there must be some point a 2 $\overline{A}$ with

$$d(q; a) \; D:$$

By denition of D, we also have

$$d(q_i; a) \; D \qquad \text{for } i = 1; 2:$$

In other words, we have a triangle $(q_1; a; q_2)$ in M where the distance from a to the midpoint q of q_1q_2 is at least both $d(a; q_1)$ and $d(a; q_2)$. Note also that because A $B(p; R^0)$, we have $q_1; q_2$ 2 $B(p; 2R^0)$, so all side lengths of our triangle are at most $4R^0$ $\frac{}{4}$:

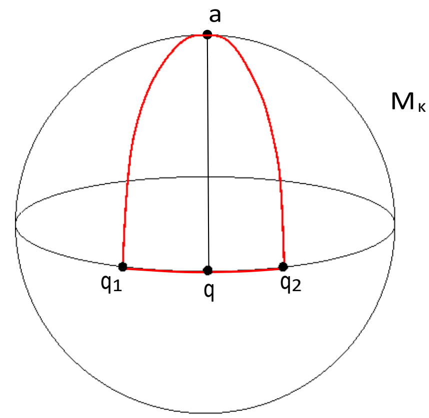

Figure 4. A triangle on a sphere in which a midpoint-vertex distance is the same as the adjacent side lengths. This happens only if the side lengths are at least half the diameter of the sphere.

We claim this is impossible. A theorem of Alexandrov, see [27, 1A.6], implies that distances between points in the triangle $(q_1; a; q_2)$ are at most those in a 'comparison triangle' with the same side lengths in the model space M with constant curvature . (Here is where we use that R^0 is less than the injectivity radius at p.) In particular, $d(q; a)$ is at most the corresponding midpoint-vertex distance of the comparison triangle. And in M, midpoint-vertex distances are always less than the maxima of the adjacent side lengths, unless M is a sphere and some side of the triangle has length at least half the diameter of the sphere, i.e. at least $\frac{p}{2}$, see Figure 4. But our triangle has side lengths at most $\frac{p}{4}$.

Claim 4.15. After possibly shrinking V, there is a well-dened, continuous map s

$$: V \;!\; FP^d; \quad s(X) = (A_X);$$

where is the circumcenter map on L_X, as in Lemma 4.14.

This will nish the proof of Lemma 4.11, for as dened above, s is constant on $_1$-equivalence classes and X $_1$ s(X). So, it remains to prove the claim.

Proof. Given > 0, we can choose V small enough such that for every X 2 V, the $_1$-equivalence class [X] L_X has d_{L_X} -diameter at most . (The argument is almost the same as that used to prove Lemma 4.8.) As injectivity radius and sectional curvature near the base point vary continuously in FM^d, we can assume that V is small enough so that by Lemma 4.14, each subset [X] L_X has a well-dened circumcenter. On such a neighborhood V, the map in the statement of the claim is well-dened.

For continuity, suppose that X_i ! X in V. Let

$$F_i : L_X \dashrightarrow L_{X_i}$$

be the almost isometric maps of (26). Combining (the proof of) Lemma 4.10 and Claim 4.13, we have

$$F_i^{\;1}(A_{X_i}) \;!\; A_X \tag{29}$$

in the Chabauty topology on closed subsets of L_X. As these sets are all contained in a compact subset of L_X, Chabauty convergence means that $F_i^{-1}(A_{X_i})$ and A_X are contained in ϵ_i-neighborhoods of each other, with $\epsilon_i \to 0$ as $i \to \infty$.

We need to show that if $p_i = \mathcal{O}(A_{X_i})$ is the circumcenter in L_{X_i}, then

$$F_i^{-1}(p_i) \to \mathcal{O}(A_X);$$

the corresponding circumcenter in L_X. After passing to a subsequence, the points $F_i^{-1}(p_i)$ converge to a point $p \in L_X$. If R_i is the minimum radius of a closed ball around p_i that contains A_{X_i}, then as the F_i are almost isometric, the sets $F_i^{-1}(B(p_i;R_i))$ Chabauty converge to a ball $B(p;R) \supset A_X$, where $R_i \to R$. But A_X cannot be contained in a ball with radius less than R, since then a slight enlargement of such a ball would contain $F_i^{-1}(A_{X_i})$ for large i, contradicting the fact that $R_i \to R$. So, p is a circumcenter for A_X.

4.3.3. Constructing the charts. Let us recall our setup. We have the two spaces

$$\mathcal{P}^d = \left\{ (M;p;D) \;\middle|\; \begin{array}{c} M \text{ a complete Riemannian } d\text{-manifold,} \\ p \in M; \text{ and } D \subset M \text{ a closed subset} \end{array} \right\} \Big/ \begin{array}{c}\text{such that } \exists \text{ an isometry } f : M \to M'; \ f(D) = D'\end{array} = \text{isometry}$$

$$\mathcal{FP}^d = \left\{ (M;e;D) \;\middle|\; \begin{array}{c} M \text{ a complete Riemannian } d\text{-manifold,} \\ e \in FM; \text{ and } D \subset FM \text{ a closed subset} \end{array} \right\} \Big/ \begin{array}{c}\text{such that } \exists \text{ an isometry } f : M \to M'; \ Df(D) = D'\end{array} = \text{isometry};$$

together with the projection map $\pi : \mathcal{FP}^d \to \mathcal{P}^d$. The relevant smooth-Chabauty topologies are discussed in §A.5. Note that π is an open map: for if $(M;p;D)$ and $(M';p';D')$ are close in $\mathcal{P}^d$, there is a (locally defined) almost isometric map f between them that takes p to p', and then given $e \in FM_p$ we have that $(M;e;D)$ and $(M';Df(e);D')$ are close in $\mathcal{FP}^d$.

Choose a point $O \in \mathcal{P}^d$, and a point $F \in \mathcal{FP}^d$ with $\pi(F) = O$. Let V be a neighborhood of F that is small enough so that Lemma 4.11 applies, and so that $U = \pi(V)$ satisfies the assumptions of Lemma 4.8. Since the equivalence relation $\sim_1$ on V is a π-pullback, Lemma 4.11 gives a continuous map

$$s : U/\sim_1 \to \mathcal{FP}^d$$

such that $\pi(s([X])) \sim_1 X$ for all $X \in U$. We define a chart h :

$$\mathbb{R}^d \times U/\sim_1 \to \mathcal{P}^d;$$

as follows. Each $Z \in \mathcal{FP}^d$ gives an exponential map

$$\exp_Z : \mathbb{R}^d \to L_{\pi(Z)} \subset \mathcal{P}^d;$$

where if $Z = (M;e;D)$ then $\exp_Z$ is the exponential map of M, with respect to the frame $e \in FM$, but composed with the leaf inclusions of Lemma 4.6, so that it can be considered as a map into $L_{\pi(Z)} \subset \mathcal{P}^d$. Then

$$h(v;[X]) := \exp_{s([X])}(v):$$

Since injectivity radius at the base point varies continuously in $\mathcal{P}^d$, after possibly shrinking V we may fix $0 < \epsilon < \frac{1}{4}$ such that for each $X \in U$, the map

$$\exp_{s([X])} : B(0;2\epsilon) \to L_{s([X])} \tag{30}$$

is a diffeomorphism onto its image. We claim:

Claim 4.16. $h : B(0;\epsilon) \times U/\sim_1 \to \mathcal{P}^d$ is a homeomorphism onto its image.

Proof. For convenience, we work with the closed ball $\overline{B(0;\epsilon)}$. We'll show that

$$h : \overline{B(0;\epsilon)} \times U/\sim_1 \;\to\; \mathcal{P}^d$$

is a continuous, proper injection. As $\mathcal{P}^d$ is Chabauty, this will imply that h is a homeomorphism onto its image.

Injectivity follows immediately from the definition of ϵ: the reason 2 appears in (30) is to ensure the exponential maps stay injective on the closed balls $\overline{B(0;\epsilon)}$.

For continuity, remember that $h(v;[X]) = \exp_{s([X])}(v)$ and note that

$$\exp : F\mathcal{P}^d \times \mathbb{R}^d \to \mathcal{P}^d, \quad (Z;v) \mapsto \exp_Z(v)$$

is continuous, since the Riemannian exponential map varies smoothly when the metric is varied smoothly, a consequence of the smooth variation of solutions to smoothly varying families of ODEs. Hence, h is continuous.

We now claim that h is proper. Assume that $(v_i;[X_i])$ is a sequence in $\overline{B(0;\epsilon)} \times U/\sim_1$, and that $h(v_i;[X_i]) \to Y \in \mathcal{P}^d$. As

$$d_{L_{X_i}}\big(X_i; h(v_i;[X_i])\big) < \epsilon,$$

Lemma 4.7 2) implies that (X_i) has a subsequence that converges to some $X \in \mathcal{P}^d$. By compactness, v_i has a subsequence that converges to some $v \in \overline{B(0;\epsilon)}$, so the sequence $(v_i;[X_i])$ is pre-compact in $\overline{B(0;\epsilon)} \times U/\sim_1$.

We now want to show that the set of all maps h constructed as above is a foliated atlas for $\mathcal{P}^d$. The key is the following lemma:

Lemma 4.17. Suppose that U and ϵ are chosen to be small enough so that $\sim_1$ is an equivalence relation on the image $h(B(0;\epsilon) \times U/\sim_1)$. Then if

$$(v;[X]);(w;[Y]) \in B(0;\epsilon) \times U/\sim_1; \text{ we}$$

have that $[X] = [Y] \iff h(v;[X]) \sim_1 h(w;[Y])$.

Proof. The forward direction is immediate, since if $s([X]) = (M;e;D)$ then $h(v;[X]) = (M;\exp_e(v);D)$ and $h(w;[Y]) = (M;\exp_e(w);D)$. Since $\epsilon < \frac{1}{4}$, we have $v,w \in B\left(0;\frac{1}{2}\right)$. So $d_M(\exp_e(v);\exp_e(w)) < 1$, as desired.

Conversely, suppose $h(v;[X]) \sim_1 h(w;[Y])$. Then

$$X \sim_1 (s([X])) \sim_1 h(v;[X]) \sim_1 h(w;[Y]) \sim_1 (s([Y])) \sim_1 Y.$$

Lemma 4.17 implies that transition maps between the charts h have the form

$$t : B \times U/\sim_1 \to B^0 \times U^0/\sim_1, \quad t(v;[X]) = (t_1(v;[X]); t_2([X])),$$

where $B;B^0$ are neighborhoods of the origin in $\mathbb{R}^d$ and $U;U^0$ are open in $\mathcal{P}^d$. Furthermore, for each fixed $[X]$, the map $v \mapsto t_1(v;[X])$ is a transition map between two exponential maps for the same Riemannian manifold M, but taken with respect to different base frames. So, t_1 is smooth in v. As X varies in $\mathcal{P}^d$, the Riemannian manifolds M vary smoothly, and the base frames vary continuously, so the v-partial derivatives of t_1 also vary continuously.

Therefore, $\mathcal{P}^d$ is a foliated space. The charts $h : B(0;\epsilon) \times U/\sim_1 \to \mathcal{P}^d$ above are chosen exactly so that when $[X_i] \to [X]$ in $U/\sim_1$, the pullback metrics g_{X_i} on $B(0;\epsilon)$ converge in the C^∞ topology to g_X. The point is just that if a sequence of framed manifolds converges in the framed smooth topology, then the almost isometric maps of Equation (45) in §A.3 can be chosen to almost commute with the exponential maps in the small neighborhoods of the base frames. Finally, by

Lemma 4.17, the leaf equivalence relation of P^d is generated by $_1$, so leaves are obtained as promised by fixing $(M;D)$ and varying the base point $p \in M$.

4.4. Application: an ergodic decomposition theorem. A unimodular probability measure μ on M^d is ergodic if whenever B is saturated and Borel, either $\mu(B) = 0$ or $\mu(B) = 1$. Here, we use the desingularization theorem (Theorem 1.8) to show that any unimodular probability measure on M^d can be expressed as an integral of ergodic such measures.

Ergodic decomposition theorems are usually proved in one of two ways. Phrased in our context, one of the usual approaches is to disintegrate μ with respect to the σ-algebra of saturated subsets, see [48], and then prove that the conditional measures are ergodic. The other approach uses that ergodic probability measures are the extreme points of the convex set of all unimodular probability measures, and then appeals to Choquet's theorem [82].

Neither of these approaches quite applies to unimodular measures on M^d in full generality. For the first approach, the usual way to prove that the conditional measures are ergodic is to appeal to a pointwise ergodic theorem. While Garnett [51] has proved a foliated ergodic theorem with respect to Brownian motion — this could be applied in our setting using Theorem 1.8 — her theory requires that the leaves of the foliation have uniformly bound geometry. The problem with the second approach is that to our knowledge, there is currently no version of Choquet's theorem that applies in this generality. Namely, the original requires that the convex set in question be compact, see [82], and more general versions such as Edgar's [43] require separability assumptions and that the underlying Banach space has the Radon-Nikodym property.

Our approach is to use Theorem 1.8 to reduce the problem to an ergodic decomposition theorem for Polish foliated spaces, and then to reduce that to a decomposition theorem for measures on a complete transversal. Essentially, if one traces through all the arguments in the referenced papers, in particular in [55], the argument does boil down to Choquet's theorem, but considering measures on the transversal allows one to use Varadarajan's compact model theorem [92, Theorem 3.2] (or rather, the easier countable case thereof) to reduce everything to the case of an ergodic decomposition for a countable group acting by continuous maps on a compact metric space, which is a setup that Choquet's theorem [82] can handle.

Proposition 4.18 (Ergodic decomposition)**.** *Let μ be a unimodular Borel probability measure on M^d. Then there is a standard probability space $(X;\nu)$ and a family $\{\mu_x \mid x \in X\}$ of ergodic unimodular Borel probability measures on M^d such that for every Borel $B \subset M^d$, the map $x \to \mu_x(B)$ is Borel and*

$$\mu(B) = \int_X \mu_x(B)\, d\nu.$$

Proof. Let P^d be the Polish foliated space defined in Theorem 1.8. We say that a completely invariant probability measure μ on P^d is ergodic if every Borel, leaf-saturated subset has measure 0 or 1. The leaf map of Proposition 4.1 pushes forward (ergodic) completely invariant measures on P^d to (ergodic) unimodular measures on M^d. In light of Theorem 1.8, it then suffices to show that for every completely invariant Borel probability measure μ on P^d, there is a standard probability space $(X;\nu)$ and a family $\{\mu_x \mid x \in X\}$ of ergodic completely invariant Borel probability measures on P^d such that for every Borel B P^d, the map $x \to {}_x(B)$ is Borel and

$$(B) = \int {}_x(B)\, d:$$

Choose a complete Borel transversal T for the foliated space P^d, i.e. a Borel subset that intersects each leaf in a nonempty countable set, and assume T is a Polish space. (The existence of such a T follows from the fact that P^d is Polish, and from the foliated structure.) The leaf equivalence relation restricts to an equivalence relation on T with countable equivalence classes, and T with its Borel -algebra is a standard Borel space. So by Feldman-Moore [45], is the orbit equivalence relation of some Borel action G T of a countable group G.

The measure is the result of integrating the Riemannian measures on the leaves of P^d against a holonomy invariant transverse measure on T. Since the action of the holonomy groupoid generates and preserves , the action G T above must also preserve , by Corollary 1 of [45].

We now apply the ergodic decomposition of Greschonig-Schmidt [55] to the transverse measure . They show that there is a standard probability space $(X;)$ and a family $\{{}_x \mid x \in X\}$ of G-ergodic Borel probability measures on T such that for every Borel B T, the map $x \to {}_x(B)$ is Borel and

$$(B) = \int {}_x(B)\, d:$$

Integrating each ${}_x$ against the Riemannian measures on the leaves of P^d gives a completely invariant measure ${}_x$ on P^d, and it follows that for every Borel B P^d, the map $x \to {}_x(B)$ is Borel and

$$(B) = \int {}_x(B)\, d:$$

5. Compactness theorems

Cheeger-Gromov's $C^{1,1}$-compactness theorem [56] states that for every $c, > 0$, the set of pointed Riemannian d-manifolds (M, p) such that

1) $|K_M()|$ c for every 2-plane , 2)
$\mathrm{inj}_M(p) > 0$

is precompact with respect to Lipschitz convergence, and that the limits are $C^{1,1}$-manifolds, Riemannian manifolds where the metric tensor is only Lipschitz. Here, K_M is the sectional curvature tensor and $\mathrm{inj}_M(p)$ is the injectivity radius of M at p. Variants of this theorem, see e.g. [81, Chapter 10], ensure greater regularity of the limits when the derivatives of K_M are bounded. For instance, if we x some sequence (c_j) in $\mathbb{R}$ and replace 1) by

1') $|r^j K_M|$ c_j, for all $j \in \mathbb{N}$ $\{0\}$,

then the space of pointed Riemannian d-manifolds (M, p) satisfying 1') and 2) is compact in the smooth topology, i.e. as a subset of M^d, see Lessa [68, Theorem 4.11]. We also discuss a similar compactness theorem in §A.1.

By the Riesz representation theorem and Alaoglu's theorem, the set of Borel probability measures on a compact metric space is weak* compact, so in particular we have weak* compactness for probability measures supported on the space of pointed manifold satisfying 1') and 2) above. As unimodularity is weak* closed, this also gives compactness for unimodular probability measures.

For the Cheeger-Gromov compactness theorems, a lower bound on injectivity radius at the base point is necessary, as no sequence $(M_i; p_i)$ with $\mathrm{inj}_{M_i}(p_i) \to 0$ can converge in the smooth (or even lipschitz) topology. In this section, however, we prove that for manifolds with pinched negative curvature, a condition on the injectivity radius is not necessary if we are only interested in the weak* compactness of the space of unimodular probability measures.

More precisely, xing $a; b > 0$ and a sequence $c_j \in \mathbb{R}$, let $\mathcal{M}^d_{PNC} \subset \mathcal{M}^d$ be the set of all pointed d-manifolds $(M; p)$ satisfying 1') and

3) $a^2 \leq K_M() \leq b^2 < 0$ for every 2-plane ,

and let $\mathcal{M}^d_{PNC;inj}$ be the subset of all $(M; p)$ that satisfy 1'), 2) and 3). The former space $\mathcal{M}^d_{PNC}$ is not compact, but even so we have:

Theorem 1.12 (Compactness in pinched negative curvature). The space of unimodular probability measures on $\mathcal{M}^d_{PNC}$ is compact in the weak topology.

The point is that the -thin part of a manifold M with pinched negative curvature only takes up a uniformly small proportion of its volume. More precisely,

Proposition 5.1 (Thick at the basepoint). If is a unimodular probability measure on $\mathcal{M}^d_{PNC}$ and > 0, there is some $C = C(; d; a; b)$ such that

$$\lim_{\to 0} C = 0,$$
$$(\mathcal{M}^d_{PNC;inj}) \geq 1 - C:$$

Since each $\mathcal{M}^d_{PNC;inj}$ is compact, Prokhorov's theorem [23, IX.65] implies that the set of unimodular probability measures on $\mathcal{M}^d_{PNC}$ is weak* precompact in the space of all probability measures, and therefore compact since unimodularity is a closed condition. So, Theorem 1.12 follows from Proposition 5.1. In fact, since any sequence $(M_i; p_i) \in \mathcal{M}^d_{PNC}$ such that $\mathrm{inj}_{M_i}(p_i) \to 0$ must diverge, Theorem 1.12 and Proposition 5.1 are formally equivalent.

We defer the proof of Proposition 5.1 to the next section, and nish here with some remarks about Theorem 1.12. First, as in the compactness theorems for pointed manifolds, control on the derivatives of sectional curvature is not necessary if one is willing to consider limits that are measures supported on $C^{1;1}$-Riemannian manifolds, and where the convergence is Lipschitz. The derivative bounds in 1') do not factor into the proof of Proposition 5.1, and are used only when appealing to the compactness of $\mathcal{M}^d_{PNC;inj}$. (See [81, Ch 10].)

Second, suppose that M_i is a sequence of nite volume Riemannian manifolds and $_i \to 0$ is a sequence such that

$$\frac{\mathrm{vol}((M_i)_{<_i})}{\mathrm{vol}\, M_i} \nrightarrow 0; \tag{31}$$

where $(M_i)_{<_i}$ is the $_i$-thin part of M_i, i.e. the set of points with injectivity radius less than $_i$. Then the corresponding unimodular probability measures $(_i = \mathrm{vol}(M_i))$, see §2, form a sequence with no convergent subsequence. So for example, a uniform lower bound on curvature is required in Theorem 1.12, since when the metric on a closed hyperbolic surface is scaled by $1=i$, the whole manifold will be $_i$-thin for some $_i \to 0$. A similar argument shows that there is no analogue of Theorem 1.12 for at manifolds.

Example 5.2. The uniform negative upper bound curvature is also necessary. To see this, construct Riemannian surfaces S_i by cutting a hyperbolic surface along a closed geodesic with length $_i$! 0, and inserting a at annulus A_i with boundary length $_i$ and width 1=$_i$ in between. The surfaces S_i have bounded volume, and the $_i$-thin part of S_i has volume at least some constant, so (31) holds. The metrics on the S_i can then be perturbed so that the metric is smooth everywhere, and slightly negative on the annuli A_i, without aecting (31).

Although there is no general analogue of Theorem 1.12 in nonpositive curvature, there is a similar compactness result for nonpositively curved locally symmetric spaces. In §5.2, we will give an algebraic proof of this, and will also indicate how our geometric arguments might be adapted to this setting. We will also discuss the possibility of a universal theorem that implies both Theorem 1.12 and its locally symmetric analogue.

5.1. The proof of Proposition 5.1. The idea is simple. One needs to show that in a manifold with pinched negative curvatures, the -thin part takes up a small proportion of the volume when is small. Then one transfers this estimate to using unimodularity. Of course, our manifolds and their thin parts may have innite volume, so one needs a local version of 'small proportion' that is robust enough to work in this setting. More precisely, we will show how to push volume from the -thin part into a region near the boundary of the $_0$-thin part, where $_0$ is the Margulis constant, without incurring a large Radon-Nikodym derivative.

Before starting the proof in earnest, we record the following facts, which should be well known to those familiar with the literature.

Lemma 5.3. Let M be a simply connected Riemannian d-manifold with curvature a^2 K_M 0. Below, all geodesics have unit speed.

1) If ; are geodesics that intersect at (0) = (0) with angle , then

$$d((t);(t)) \quad \frac{\sinh(at)}{a}:$$

2) If ; are geodesics that both intersect a geodesic orthogonally at (0) and (0), and $\frac{d}{dt}(0);\frac{d}{dt}(0)$ are parallel vectors along , then

$$d((t);(t)) \quad d((0);(0)) \cosh(at):$$

3) If 2 $@_1$M and ; are geodesics such that (0) and (0) lie on a common horosphere around and $\lim_{t!1}(t) = \lim_{t!1}(t) =$,

$$d((t);(t)) \quad d((0);(0))\, e^{at};$$

4) Given ;T > 0, there is some = (;T;a) > 0 such that if ; are geodesics and d((t);(t)) for all t 2 [0;1], then

$$d((t);(t)) \;;\quad 8t\ 2\ [0;T]:$$

Proof. 1) is an immediate consequence of Toponogov's theorem. 2) follows from Berger's extension of Rauch's comparison theorem [39, Theorem 1.34], by interpolating between and by a one parameter family of geodesics $_s(t)$ such that $\frac{d_s}{dt}j_{t=0}$ is a parallel vector eld along :

For 3), let be a triangle in the model space $H^2_{a^2}$ with curvature a^2, chosen with one ideal vertex and so that the other two vertices lie on a common horocycle

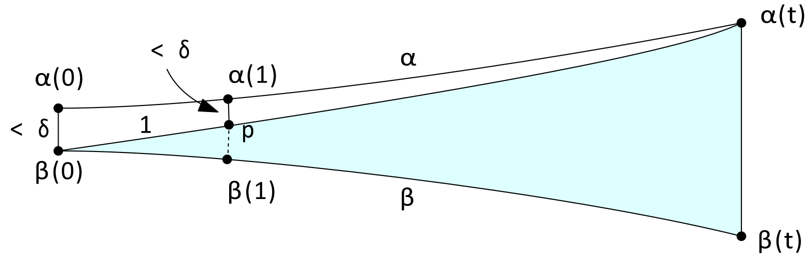

Figure 5. Above, d((1); p) < by convexity of the distance function. Then d((1); p) < 2, so applying Toponogov's theorem to the blue triangle gives an explicit upper bound for d((t); (t)) that depends only on ; t and a. This upper bound is that which one gets in $H^2_{a^2}$, so goes to zero as ! 0. So for small , we have d((t); (t)) for t T.

centered at , at a distance of d((t); (t)) from each other. Parametrize the innite sides of by arc length using ; : (1; t] ! H_{a^2}. By Toponogov's theorem[8], d((0); (0)) d((0); (0)). If is the geodesic in H_{a^2} from (0) to (0), we can ow every point on away from a distance of t to produce a path $_t$ from (t) to (t). In the upper half plane model for $H^2_{a^2}$ with = 1, the path $_t$ is created by dividing all the y-coordinates of the points on by e^{at}. Hence, length($_t$) = e^{at} length(). This gives the estimate

$$d((t); (t)) = d((t); (t)) \quad d((0); (0))e^{at} \quad d((0); (0))e^{at}:$$

For 4), see Figure 5.

Proof of Proposition 5.1. Let $_0$ be the Margulis constant for manifolds M with curvature bounds a^2 K_M () b^2 < 0, and let M̃ be the universal cover of M. If 0 < $_0$, and M is such a manifold, consider the set $M_<$ of all points p 2 M with inj_M (p) < . By the Margulis lemma, see [12, §10], the components N of $M_<$ come in two types.

1) Margulis tubes. N is the quotient of a tubular neighborhood Ñ M̃ of a geodesic ~ in M̃ by an innite cyclic group of hyperbolic-type isometries that stabilize . Hence, N is a tubular neighborhood of a closed geodesic in M, which we call its core geodesic.
2) Cusp neighborhoods. N is the quotient of an open set Ñ M̃ by a virtually nilpotent group of parabolic isometries that all have a common xed point 2 $@_1$M̃.

In both cases, the closure of N is a codimension-0 submanifold of M that has piecewise-smooth boundary. To see this, given any compact set K M̃, proper discontinuity implies that Ñ \ K is a nite union of sets of the form

$$U_g \setminus K; \text{ where } U_g = fx\ 2\ \tilde{M}\ j\ d(x; g(x)) < g;$$

[8] This ideal version of Toponogov's theorem follows the usual one since is a limit of comparison triangles associated to (t); (t); (T), as T ! 1: For T << 0, these triangles will be almost isosceles (since (t); (t) lie on a common horosphere) which implies that the limit triangle in $H^2_{a^2}$ will have two vertices on a horocycle centered at the other vertex.

and the g are deck transformations. So, it suces to show that the frontier $@U_g$ of each U_g is a smooth submanifold. But $@U_g$ is cut out by the equation $d(x; g(x)) =$, and we claim that is a regular value for $x !$ $d(x; g(x))$.

As $\tilde{M}$ is a simply connected manifold with negative curvature, it is uniquely geodesic, so the distance function of $\tilde{M}$ is smooth o the diagonal. But g has no xed points, so this means that the map $x !$ $d(x; g(x))$ is smooth. Suppose that $x 2 \tilde{M}$ is a critical point for this map, and let : $[0; D]$! $\tilde{M}$ be the geodesic from x to $g(x)$. As we start to move along from x, the distance $d(x; g(x))$ is constant to rst order. So, by the rst variational formula, $dg(^0(0)) = {}^0(D)$. Hence, the biinnite geodesic extending is invariant under g. This would mean that is the axis of minimal displacement for g, which is impossible since g translates points on by , but has translation length on M less than .

Let $M_<$ be the subset obtained from $M_<$ by removing the core geodesics of all Margulis tubes, let $@M_<$ be the boundary of the closure of $M_<$, and dene $M_{<_0}$ and $@M_{<_0}$ similarly. Note that there is a natural foliation of $M_{<_0}$ by (the interiors of) geodesics with one endpoint on $@M_{<_0}$, and where the other end either terminates at an orthogonal intersection with the core of a Margulis tube, or is a geodesic ray exiting a cusp[9]. We will call these geodesic leaves.

Claim 5.4. There is some $C = C(; d; a)$, with $C ! 1$ as $! 0$, such that the distance along any geodesic leaf from $@M_<$ to $@M_{<_0}$ is at least C.

Proof. Fixing some $R > 0$, we want to show that if is suciently small, then the distance along any geodesic leaf from $@M_<$ to $@M_{<_0}$ is at least R.

This is easiest to show for components N $@M_<$ that are neighborhoods of cusps. If $\tilde{N}$ $\tilde{M}$ is a component of the preimage of N, the geodesic leaves of N lift to geodesic rays in $\tilde{N}$ that limit to a point $2 @_1\tilde{M}$. If $\tilde{x} 2 \tilde{N}$, there is some deck transformation g that is parabolic with $g() =$ and $d(x; g(x)) <$:

Let be the unit speed geodesic in $\tilde{M}$ with $(0) = x$ and $\lim_{t! 1} (t) =$: The image $= g$ is also a geodesic limiting to , and for all t, the points (t) and (t) lie on a common horosphere centered at . By Lemma 5.3 3),

$$d((t); (t)) \quad a : \frac{e^{at}}{} \tag{32}$$

So as long as $t < C := \frac{1}{a} \log(\frac{a_0}{})$, the element g will move (t) at most $_0$. So, the distance along from $@M_<$ to $@M_{<_0}$ is at least C.

Now consider a component of $M_{<_0}$ that is a Margulis tube with core geodesic c, which we lift to a neighborhood $\tilde{N}$ of a geodesic $\tilde{c}$ in $\tilde{M}$. Let be the group of deck transformations stabilizing $\tilde{N}$. Here, all the non-identity elements act as nontrivial translations along $\tilde{c}$, coupled with the action of some element of $O(d$ $1)$ in the orthogonal direction. We claim

> (?) Given $C > 0$, there is some > 0 such that when the length of the core geodesic c is less than , we have $d(c; @M_{<_0})$ C:

[9] In the universal cover $\tilde{M}$, the normal exponential map exp : N ! $\tilde{M}$ is a dieomorphism for any geodesic . For a given g, the displacement function $x !$ $d(x; g(x))$ is convex along (the orthogonal) geodesics, so each U_g from before is a star-shaped neighborhood of . The picture is similar for cusp neighborhoods, except that now we consider the foliation of $\tilde{M}$ by geodesics with one endpoint at some $p 2 @_1\tilde{M}$.

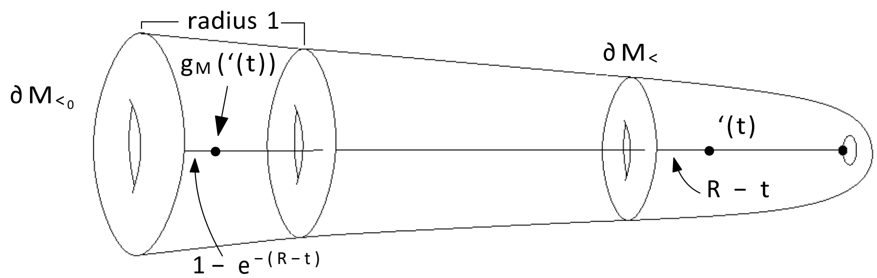

Figure 6. The denition of the map g_M , pictured in a Margulis tube in dimension 3. Note that the image of g_M lies in a radius 1 neighborhood of $@M_{<_0}$:

To prove this, rst use Lemma 5.3 1) to choose some > 0 small enough (= a_0=(2 sinh(a)) works) so that any two geodesic rays $_1$(t); $_2$(t) in M that intersect at t = 0 with an angle at most satisfy

(33) $$d(_1(t);\ _2(t))\ \ _0=2; \quad \text{for all } t\ \ C:$$

Since O(d 1) is compact, there is some n 2 N such that any A 2 O(d 1) has a power A^k with k = k(A) n that is close enough to the identity so that

$$\backslash(A^k(v); v) < \ ; \quad \text{for all } v\ 2\ R^{d\ \ 1}:$$

Finally, using Lemma 5.3 2), let > 0 be small enough so that any two geodesic rays $_1$(t); $_2$(t) in H^d that start out perpendicular to a third geodesic, parallel to each other and at distance at most from each other, must also satisfy (33). (For instance, take = $\frac{0}{2\cosh(a)}$:) Now take = $\frac{}{n}$:

If g 2 is any isometry with translational part less than , then for some k n, the isometry g^k has translational part at most and rotates vectors orthogonal to the core geodesic c by at most an angle of . From this, we see that g^k has displacement at most $_0$ everywhere on the C-neighborhood of c. So in any component of $M_{<_0}$, the distance from the core curve c to the boundary $@M_{<_0}$ is at least C whenever the core has length less than , proving (?).

To complete the proof of the claim, given C > 0 we choose as in (?), but using C + 1 instead of C. Using Lemma 5.3 4), we may also assume that is small enough so that any two geodesics $_1$(t); $_2$(t) in M satisfy

(34) $$d(_1(t);\ _2(t))\ \ ;\ 8t\ 2\ [0;1] \ =)\ \ d(_1(t);\ _2(t))\ \ _0;\ 8t\ 2\ [0; C+1]:$$

We want to show that the distance from $@M_<$ to $@M_{<_0}$ is at least C. If the distance from the core c to $@M_<$ is less than 1, then we are done since we know by (?) that d(c; $@M_{<_0}$) C + 1. So, we may assume that the distance from the core to $@M_<$ is at least 1.

If ‘ is a geodesic leaf in $M_{<_0}$, parameterized by arc length so that ‘(1) is on the boundary $@M_<$, there are two lifts $_1$; $_2$ of ‘ in N M with d($_1$(t); $_2$(t)) throughout the preimage of $M_<$, so for all t 2 [0; 1]. Therefore, (34) implies that the injectivity radius along ‘ stays less than $_0$ for all t 2 [0; C + 1], i.e. for at least a length of C after exiting $M_<$.

We now dene a map

$$g_M : M_< \ \ ! \ \ M_{<_0}$$

as follows. If ℓ is a geodesic leaf, parameterized by arc length so that $\ell(0)$ is on the core, suppose that $\ell(R) \in \partial M_{<\epsilon}$ and $\ell(R_0) \in \partial M_{<\epsilon_0}$. Then define

$$g_M(\ell(t)) = \ell(R_0 - 1 + e^{b(t-R)}).$$

In other words, g_M is constructed to map $M_{<\epsilon}$ to a 1-neighborhood of $\partial M_{<\epsilon_0}$, as shown in Figure 6. This g_M is a piecewise smooth homeomorphism onto its image, since R and R_0 are piecewise smooth functions of the point $\ell(t) \in M_{<\epsilon}$.

Claim 5.5. At every point p where g_M is smooth, we have

$$|\det dg_M(p)| \geq D(\epsilon),$$

for some $D(\epsilon)$ that tends to infinity as $\epsilon \to 0$.

Here, the determinant is calculated with respect to orthonormal bases in the two relevant tangent spaces, and so measures the volume distortion of g_M.

Proof. First, note that at $p = \ell(t_0)$, the map g_M stretches lengths in the direction of ℓ by a factor of $\frac{d}{dt}(R_0 - 1 + e^{b(t-R)})|_{t=t_0} = be^{b(t_0-R)}$.

Assume now that we are working within a component of $\partial M_{<\epsilon_0}$ that is a Margulis tube with core geodesic c. Given a point p at which g_M is smooth, let $\ell_s(t)$ be a one parameter family of geodesic leaves satisfying:

1) $\ell_0(t_0) = p$ and $|\frac{d}{ds}\ell_s(t_0)| = 1$.
2) $|\frac{d}{dt}\ell_s(t)| = 1$ for all s, t, and $\ell_s(0) \in c$ for all s,

The path $s \to \ell_s(t_0)$ passes through p at $s = 0$ and moves along the boundary $\partial B(c;t)$ of the radius t_0-neighborhood of c in M, orthogonally to the leaves. Its image under g_M is just the path $s \to \ell_s(R_0 - 1 + e^{a(t_0-R)})$.

The equation $J(t) = \frac{d}{ds}\ell_s(t)|_{s=0}$ defines a Jacobi field along ℓ_0, and

$$dg_M(J(t)) = J(R_0 - 1 + e^{b(t-R)}).$$

As the path $s \to \ell_s(t_0)$ lies in $\partial B(c;t_0)$, and $B(c;t_0) \subset M$ is convex, Warner's extension of the Rauch comparison theorem [94, Theorem 4.3 (b)] implies that for $t \geq t_0$ we have[10] $|J(t)| \geq \cosh(b(t-t_0))$. So in particular,

$$\begin{aligned} |J(R_0 - 1 + e^{b(t_0-R)})| &\geq \cosh(b(R_0 - 1 + e^{b(t_0-R)} - t_0)) \\ &\geq \frac{1}{2}e^{b(R_0-1-t_0)} \\ &= \frac{1}{2}e^{b((R_0-R)-1)}e^{b(R-t_0)}. \end{aligned}$$

[10] Here, $\frac{\cosh(b(t-t_0))}{b}$ is the length of a Jacobi field in $\mathbb{H}^2_{-b^2}$ obtained by differentiating a unit speed geodesic variation $\gamma_s(t)$, where $s \to \gamma_s(t_0)$ is also a unit speed geodesic, and is perpendicular to all the geodesics $t \to \gamma_s(t)$. Warner's theorem requires the sectional curvatures in M to be less than those in the comparison space, the two Jacobi fields to have the same length at $t = t_0$, and the path $s \to \ell_s(t_0)$ to lie in a codimension one submanifold S orthogonal to $\frac{d}{dt}\ell_s(t)$ all of whose principal curvatures are larger than those of a corresponding submanifold S' containing $s \to \gamma_s(t_0)$. In our case, both Jacobi fields have length one at $t = t_0$. Convexity implies that the principal curvatures of $S = \partial B(c;t_0)$ are nonnegative, when calculated with respect to the outward normal $\frac{d}{dt}\ell_s(t)|_{t=t_0}$, so they are larger than those of the geodesic $s \to \gamma_s(t_0)$, which are zero.

As $|J(t_0)| = 1$, this means that dg_M scales the length of $J(t_0)$ by at least a factor of $\frac{1}{2}e^{b((R_0-R)-1)}e^{b(R-t_0)}$. But above, $J(t_0)$ can be taken to be any vector in TM_p orthogonal to ℓ, by choosing the variation ℓ_s appropriately. So,

$$|\det dg_M(p)| \geq \left(be^{b(t_0-R)}\,\frac{1}{2}e^{b((R_0-R)-1)}e^{b(R-t_0)}\right)^{d-1} \geq \frac{b}{2^{d-1}}e^{b((R_0-R)-1)} := D(\epsilon),$$

a constant that tends to infinity as $\epsilon \to 0$, by Claim 5.4. The argument for a component of $\partial M_{<\epsilon_0}$ that is a neighborhood of a cusp is almost exactly the same. Instead of parameterizing the geodesic variations ℓ_s so that for constant t_0, the path $s \mapsto \ell_s(t_0)$ lies along a metric sphere around the core of the Margulis tube, we parameterize so that $s \mapsto \ell_s(t_0)$ is contained in a horosphere. Horospheres are C^2 [61] and convex [42], so one can still apply Warner's comparison theorem.

We now use unimodularity to finish the proof of Proposition 5.1. Let $\mathcal{M}^d_{PNC,2}$ be the space of all isometry classes of doubly pointed d-manifolds (M,p,q) with pinched negative curvature $-a^2 \leq K_M \leq -b^2 < 0$ and geometry bounds as in 1') at the beginning of the section. Define a Borel function $F : \mathcal{M}^d_{PNC,2} \to \mathbb{R}_+$ via

$$F(M,p,q) = \begin{cases} |\det dg_M(p)| & \text{if } p \in M_{<\epsilon} \text{ and } d(g_M(p),q) \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

Using Claim 5.4, fix some $\epsilon < \epsilon_0$ with $C(\epsilon) \geq 1$. By definition, the image of g_M lies in a radius 1 neighborhood of $\partial M_{<\epsilon_0}$, so the injectivity radius at every $g_M(p)$ is at least ϵ. Setting $V(\rho,\kappa)$ to be the volume of a radius ρ ball in the d-dimensional model space of constant curvature κ, we have that the volume of a ball of radius 1 around each $g_M(p)$ satisfies

$$V(\epsilon, -b^2) \leq \mathrm{vol}(B_M(g_M(p),1)) \leq V(1,-a^2).$$

So, using this and (5.5), we compute:

$$\begin{aligned}
\mu(\mathcal{H}^d_{<\epsilon}) &\leq \int_{(M,p)\in\mathcal{M}^d_{PNC,inj<\epsilon}} \frac{|\det dg_M(p)|}{D(\epsilon)}\,\frac{\mathrm{vol}(B_M(g_M(p),1))}{V(\epsilon,-b^2)}\,d\mu \\
&= \frac{1}{D(\epsilon)V(\epsilon,-b^2)} \int_{(M,p)\in\mathcal{M}^d_{PNC,inj<\epsilon}} \int_{q\in M} F(M,p,q)\,d\mathrm{vol}\,d\mu \\
&= \frac{1}{D(\epsilon)V(\epsilon,-b^2)} \int_{(M,p)\in\mathcal{M}^d_{PNC,inj<\epsilon}} \int_{q\in M} F(M,q,p)\,d\mathrm{vol}\,d\mu \qquad (35)\\
&= \frac{1}{D(\epsilon)V(\epsilon,-b^2)} \int_{(M,p)\in\mathcal{M}^d_{PNC,inj<\epsilon}} \int_{q\in M_{<\epsilon}\cap g_M^{-1}(B_M(p,1))} |\det dg_M(q)|\,d\mathrm{vol}\,d\mu \qquad (36)\\
&\leq \frac{1}{D(\epsilon)V(\epsilon,-b^2)} \int_{(M,p)\in\mathcal{M}^d_{PNC,inj<\epsilon}} \mathrm{vol}(B_M(p,1))\,d\mu \leq \frac{V(1,-a^2)}{D(\epsilon)V(\epsilon,-b^2)};
\end{aligned}$$

where here $\mathcal{M}^d_{PNC;\,inj<\epsilon}$ is the set of all $(M,p) \in \mathcal{M}^d_{PNC}$ where $inj_M(p) < \epsilon$. Equation (35) is unimodularity, while (36) is just the change of variables formula. The last line goes to zero as $\epsilon \to 0$, which proves Proposition 5.1.

5.2. Locally symmetric spaces. Let X be a symmetric space of nonpositive curvature with no Euclidean factors, and let $G = \mathrm{Isom}(X)$. An X-manifold is a quotient X/Γ, where $\Gamma < G$ is discrete and torsion free. Let $\mathcal{M}^X \subset \mathcal{M}^d$ be the subset of pointed X-manifolds.

Theorem 1.13 (Compactness for locally symmetric spaces). The space of unimodular probability measures on $\mathcal{M}^X$ is weak*-compact.

By Proposition 2.9, there is a dictionary between unimodular measures on $\mathcal{M}^X$ and discrete, torsion free invariant random subgroups of G. The space of all invariant random subgroups of G is compact, since the Chabauty topology on the set of closed subgroups of any locally compact group is compact. As G is semi-simple, c.f. [62], it suces to prove the following proposition:

Proposition 5.6. For IRSs in a semi-simple Lie group G, 'discrete' is a closed condition. For discrete subgroups $H < G$, 'torsion-free' is a closed condition.

Note that unlike the second sentence, the rst is true only on the level of IRSs, since there are always discrete, cyclic subgroups of G that limit in the Chabauty topology to subgroups isomorphic to $\mathbb{R}$.

Lemma 5.7. If G is a semi-simple Lie group, every connected IRS $H \subset G$ is a normal subgroup of the identity component $G^\circ \subset G$.

Proof. If $H < G$ is any ergodic connected IRS, the Lie algebra $\mathfrak{h}$ is a random k-dimensional subspace of $\mathfrak{g}$, for some k, and the distribution of $\mathfrak{h}$ is invariant under the adjoint action of G on the Grassmannian $Gr(k,\mathfrak{g})$. Since G has no nontrivial compact quotients, applying the arguments of [49, Lemmas 2,3] to $Gr(k,\mathfrak{g})$ instead of the projective space $P(\mathfrak{g})$, we see that the distribution of $\mathfrak{h}$ is concentrated on $\mathrm{Ad}\,G$-invariant subspaces of $\mathfrak{g}$.

Proof of Proposition 5.6. Any subgroup $H \subset G$ that is a Chabauty limit of discrete groups has a nilpotent identity component, c.f. [90, Theorem 4.1.7]. So, any IRS that is a limit of discrete IRSs also as this property. However, a semi-simple Lie group does not have any nontrivial nilpotent normal subgroups, so Lemma 5.7 implies that a limit of discrete IRSs is discrete.

Next, we show that 'torsion-free' is a closed condition within the space of discrete $H < G$. For suppose $H_n \to H$ are discrete and the only torsion is in the limit, and take $1 \neq \gamma \in H$ with $\gamma^k = 1$. Picking $\gamma_n \in H_n$ that converge to γ, the sequence (γ_n^k) consists of nontrivial elements converging to 1. Aiming for a contradiction, choose open balls $B_1, B_2, \ldots \ni 1$ in G with $B_i \supset \overline{B_{i+1}}$. Exponentiating the γ_n^k by appropriate powers determined using exponential coordinates, it follows that for any xed i, and large enough n, there is an element of H_n inside $B_i \smallsetminus \overline{B_{i+1}}$. This contradicts discreteness of H.

Using the same arguments as in the paragraph after the statement of Proposition 5.1, Theorem 1.13 is equivalent to the following:

Proposition 5.8. If μ is a unimodular probability measure on $\mathcal{M}^X$ and $\epsilon > 0$, there is some $C = C(\epsilon, X)$ such that

$\lim_{\epsilon \to 0} C_\epsilon = 0$,
$\mu(M^X_{inj \geq \epsilon}) \geq 1 - C_\epsilon$,

where $M^X_{inj \geq \epsilon}$ be the set of pointed X-manifolds (M,p) such that $inj_M(p) \geq \epsilon$.

While the algebraic proof of Theorem 1.13 is quite direct, it is natural to ask whether there is a geometric proof of Proposition 5.8. In particular, why is it true for locally symmetric spaces when it fails for general spaces of nonpositive, or even unpinched negative, curvature? (See Example 5.2.)

The biggest difference is that when X is a non-positively curved symmetric space without Euclidean factors, its Ricci curvature is bounded away from zero, and hence the same is true for any X-manifold. (In contrast, the surfaces of Example 5.2 have points where the Ricci curvature is arbitrarily close to zero.) Specifically, the Ricci curvature tensor of $X = G/K$ is a constant negative multiple of the Killing form on $\mathfrak{g}$ [11, Theorem B.24], so is negative definite. There is then some constant $b = b(X) > 0$ such that

$$\mathrm{Ric}(v,v) \leq -b^2/(d-1), \quad \forall v \in T^1X, \tag{37}$$

since X is homogenous. Here, d is dimension, and we prefer to write the upper bound as above since then it follows that every v is contained in a 2-plane with sectional curvature at most $-b^2$. We ask:

Question 5.9. Consider the space $M^d(a,b,c_j)$ of all $(M,p) \in M^d$ with

1') $|\nabla^j K_M| \leq c_j$ for all $j \in \mathbb{N} \cup \{0\}$;
3') $-a^2 \leq K_M(\sigma) \leq 0$, and $\mathrm{Ric}(v,v) \leq -b^2$, for every 2-plane σ and unit vector $v \in TM_p$,

where Ric is the Ricci curvature. Is the set of unimodular probability measures on $M^d(a,b,c_j)$ weak* compact? More concretely, does condition 3') imply a uniform bound $C = C(\epsilon, a, b, d)$ on the proportion of volume that the ϵ-thin part occupies in (say) a finite volume $(M,p) \in M(a,b,c_j)$, with $C \to 0$ as $\epsilon \to 0$?

A proof of such a compactness result would be possible, for instance, if one could prove (say, for Hadamard manifolds) that an upper bound on Ricci curvature gives volume comparison results analogous to those given by Bishop-Gromov, see [81, Lemma 7.1.4], for Ricci lower bounds.

We do expect that our proof in pinched negative curvature can be adapted to the locally symmetric setting—that is, to give a geometric proof of Proposition 5.8—without proving the volume comparison results mentioned above. Much of the argument goes through unchanged, although it becomes more subtle to develop an analogue of the foliation by geodesic leaves when the curvature is only nonpositive.

6. Unimodular random hyperbolic manifolds

Recall that the limit set of a subgroup Γ of $\mathrm{Isom}(\mathbb{H}^d)$ is the subset of $\partial \mathbb{H}_d = S^{d-1}$ consisting of all accumulation points at infinity of any orbit of $\Gamma \curvearrowright \mathbb{H}^d$. In [3], it is shown that if μ is an invariant random subgroup of $\mathrm{Isom}(\mathbb{H}^d)$ without an atom at the identity, then the limit set of $H \subset \mathrm{Isom}(\mathbb{H}^d)$ must have full limit set μ-almost surely. In two dimensions, this has the following corollary:

Corollary 6.1. Any unimodular random hyperbolic surface with finitely generated fundamental group is either $\mathbb{H}^2$ or has finite volume.

As usual, the translation between IRSs and URHSs is through Proposition 2.9. Recall that a unimodular random manifold is a random element of $\mathcal{M}^d$ whose law is a unimodular probability measure, and is hyperbolic if its law is supported on hyperbolic manifolds. For those allergic to probabilistic language, the statement of the corollary is that if μ is a unimodular measure on $\mathcal{H}^2$ then μ-a.e. pointed hyperbolic surface (S,p) with finitely generated π_1 is either isometric to $\mathbb{H}^2$ or has finite volume.

There is an alternative way to prove Corollary 6.1, using the No-Core Principle (Theorem 1.15). Recall that the convex core of a hyperbolic surface S is the smallest convex subsurface whose inclusion is a homotopy equivalence. When S has finitely generated fundamental group, its convex core is compact. Alternatively, the ends of S are geometrically either infinite volume flares or finite volume cusps. Here, a flare is cut off by a closed geodesic, and is isometric to half of the quotient of $\mathbb{H}^2$ by the group generated by a single isometry of hyperbolic type. The convex core of S is obtained by chopping off all flares at the bounding closed geodesics. See [14] for details.

Define a function $f : \mathcal{H}^2 \longrightarrow \{0,1\}$ by setting $f(S,p) = 1$ whenever p is in the open radius 1 neighborhood of the convex core of S. We claim $f^{-1}(1)$ is open, so that f is Borel. If $f(S,p) = 1$, then p lies at distance less than 1 from some closed geodesic γ in S, by density of closed geodesics in the convex core. Using the almost isometric maps defining the smooth topology, see §A.1, γ can be transferred to a closed $(1+\epsilon)$-quasigeodesic in any nearby (S',p') and then tightened to a closed geodesic γ' using the Morse Lemma [27, Theorem 1.7, pg 401]. This geodesic γ' is contained in the convex core of S', and will lie at a distance less than 1 from p as long as $(S',p') \approx (S,p)$.

So, by Theorem 1.15, we have for μ-a.e. (S,p) that

$$0 < \mathrm{vol}_S\{q \in M \mid f(S,q) = 1\} < \infty \implies \mathrm{vol}(S) < \infty.$$

As a hyperbolic surface with finitely generated π_1 either is $\mathbb{H}^2$ or has a finite volume core, this proves Corollary 6.1.

As explained in the introduction, there are examples of unimodular random hyperbolic 3-manifolds with finitely generated fundamental group other than $\mathbb{H}^3$ and finite volume manifolds, e.g. cyclic covers $\hat{M}$ of closed hyperbolic 3-manifolds M fibering over the circle. We now show that in 3-dimensions, every URHM with finitely generated π_1 looks coarsely like such $\hat{M}$. To do this, though, we need to recall some background about the geometry of ends.

Suppose that M is a hyperbolic 3-manifold with finitely generated fundamental group. Then M is homeomorphic to the interior of a compact 3-manifold, by the Tameness Theorem of Agol [7] and Calegari-Gabai [29]. Each end E of M then has a neighborhood that is a topological product $S \times (0,\infty)$, for some closed surface S, and can be classified geometrically according to its relationship with the convex core of M. Here, the convex core of M is the smallest convex submanifold of M whose inclusion is a homotopy equivalence. When M has no cusps[11], work of Bonahon [22] and Canary [31] implies that either E has a neighborhood disjoint from the convex core, in which case we call it a convex cocompact end, or it has a neighborhood

[11] In general, cusps may split the topological ends of M into 'geometric ends' with smaller genus, which have a similar classification. See [76].

completely contained in the convex core, in which case E is degenerate. See [66] for more details.

A cyclic cover $\hat{M}$ of a mapping torus is homeomorphic to S R, and both of its two ends are degenerate|the convex core of $\hat{M}$ is the entire manifold. In ABBGNRS [3, §12.5], we constructed more general examples of IRSs that give doubly degenerate unimodular random hyperbolic structures on S R; for instance, our examples do not cover any nite volume 3-manifold.

Theorem 1.14. Every unimodular random hyperbolic 3-manifold with nitely generated fundamental group either is isometric to H^3, has nite volume, or is a doubly degenerate hyperbolic structure on S R for some nite type surface S.

The proof is another application of the No-Core Principle, and the idea is that any innite volume hyperbolic 3-manifold that is not H^3 and is not a doubly degenerate hyperbolic structure on S R has a geometrically dened 'core'. Intuitively, the core should just be obtained by cutting o the ends of the manifold, but the dif-cult part is doing this in a canonical enough way that for as (M; p) varies through H^3, the condition that p lies in the core is Borel. For simplicity, we'll rst prove Theorem 1.14 only for manifolds with no cusps, and then at the end we will make some brief comments about the modications needed to extend to the general case.

Fix once and for all some > 0 less than the Margulis constant and let M be a hyperbolic 3-manifold with no cusps. The -electric distance between two points p; q in M is the inmum over all smooth paths joining p; q of the length of the intersection of with M, the -thick part of M. Given R > 0, an R-core for M is a compact 3-dimensional submanifold N M such that

1) the diameter of N is less than R,
2) N is contained in an open radius 1 neighborhood of the convex core of M,
3) the component E M r N facing each component S @N is a neighborhood of an end that is homeomorphic to S R,
4) if E, as in 3), is a neighborhood of a convex cocompact end of M, then E lies completely outside the convex core of M,
5) if E, as in 3), is a neighborhood of a degenerate end of M and p 2 E has -electric distance more than R from N, there is a level surface in E = S R that passes through p and has -electric diameter less than R.

Here, a level surface of S R is any embedded surface isotopic to a ber S ftg. The point of the denition above is that N is a small-diameter 'core' for M, obtained topologically by chopping o the ends of M. Conditions 2), 4) and 5) require that convex cocompact ends are chopped o near the convex core boundary, and the removed neighborhoods of degenerate ends have cross-sections with small electric diameter.

Every hyperbolic 3-manifold with nitely generated $_1$ and no cusps, that is not isometric to H^3, has an R-core for some R > 0. Namely, the Tameness Theorem gives such an N satisfying 1) and 3) for some R. The complement of the convex core of M consists of product neighborhoods of the convex cocompact ends [30, II.1.3], so we may pick N to satisfy 2) and 4). Finally, Canary's Filling Theorem [32] states that there is a neighborhood E of each degenerate end of M such that E is homeomorphic to S R and is exhausted by the images of simplicial hyperbolic surfaces[12] in the homotopy class of the ber, on which no essential simple closed curves of length less than are null-homotopic in M. By the Bounded Diameter Lemma [32, Lemma 4.5], such surfaces have -electric diameter bounded above by some constant depending only on and (S). So increasing R and enlarging N, we have 5) using Freedman-Hass-Scott [47] to replace the simplicial hyperbolic surfaces by embedded level surfaces, compare with [20, Corollary 3.5].

Proposition 6.2 (Borel-parametrized cores). Suppose $M = H^3$ is a hyperbolic 3-manifold with nitely generated fundamental group and no cusps, let $R > 0$ and $C_R(M)$ be the union of all R-cores of M. Then:

1) Unless M is a doubly degenerate hyperbolic metric on S R, for some closed surface S, the subset $C_R(M)$ M has nite, nonzero volume for suciently large R.
2) If we set $C_R(M) = ;$ for all other $(M; p)$ 2 H^3, the subset C_R H^3 consisting of all $(M; p)$ with p 2 $C_R(M)$ is Borel.

Both conditions 4) and 5) in the denition of an R-core are necessary for the 'nite volume' part of this proposition. Condition 4) is needed in order to prevent a sequence of R-cores from exiting the degenerate end of a hyperbolic 3-manifold homeomorphic to S R that has one convex cocompact end and one degenerate end. Less obviously, condition 5) is needed to prevent a sequence of R-cores from exiting a hyperbolic 3-manifold homeomorphic to the interior of a handlebody whose single end is degenerate. Here, the point is to express the interior of the handlebody as S R, where S is a closed, orientable surface with a single puncture, and then take appropriate 'cores' of the form $\hat{S}$ $[t$ $1; t + 1]$, where $\hat{S}$ is obtained from S by truncating its cusp.

Deferring the proof, let's understand how Proposition 6.2 implies Theorem 1.14 in the no-cusp case. Suppose is a unimodular measure on H^3, and apply the No-Core Principle (Theorem 1.15) to the characteristic function of each C_R. We obtain that for each $R > 0$, the following holds for -a.e. $(M; p)$:

$$0 < \operatorname{vol} C_R(M) < 1 \quad =) \quad \operatorname{vol} M < 1 :$$

Taking a countable union of measure zero sets, we have for -a.e. $(M; p)$ that 0

$$< \operatorname{vol} C_R(M) < 1 \text{ for some } R \ 2 \ N \quad =) \quad \operatorname{vol} M < 1 :$$

So by Proposition 6.2, we have that -a.e. $(M; p)$ with nitely generated fundamental group and no cusps is either nite volume (i.e. closed), H^3 or a doubly degenerate hyperbolic 3-manifold homeomorphic to S R.

6.0.1. Proof of Proposition 6.2. We rst prove 1), so assume that M has nitely generated $_1$, no cusps and is not isometric to H^3. As mentioned above, M admits an R-core for some R. So, for large R the subset $C_R(M)$ has nonzero volume. Our goal is to show that $C_R(M)$ is always bounded in M, so always has nite volume.

Suppose on the contrary that there is a sequence C_i of R-cores of M that exits an end E of M. By condition 2) in the denition, R-cores lie near the convex core, so E is a degenerate end of M. Increasing R if necessary, pick a neighborhood E of E that is homeomorphic to a product S R and is exhausted by level surfaces with -electric diameter at most R.

[12] A simplicial hyperbolic surface is a map from a triangulated surface that is totally geodesic on each triangle, and where the total angle around each vertex is at least 2, [32].

We claim that M is a doubly degenerate hyperbolic structure on S R. For large i, the -electric distance from C_i E to the frontier of E is at least 6R. This electric distance is realized by a path in some component D M r C_i. Then D = S^0R for some closed surface S^0, and since is contained in the convex core of M, this D cannot be a neighborhood of a convex cocompact end by condition 4) in the denition of an R-core. So, by condition 5), there are surfaces $^0_1$; $^0_2$ M that are level surfaces of D = S^0R, have -electric diameter less than R, and which pass through points on at -electric distance R and 5R, respectively, from C_i. These surfaces must then be disjoint, so they bound a submanifold N homeomorphic to S^0 [0; 1]. However, we can also pick a level surface in E = S R with -electric diameter less than R that passes through a point on at -electric distance 3R from C_i. This N = S^0 [0; 1], and is incompressible since it is incompressible in E, which contains N. Therefore, by Waldhausen's Cobordism Theorem [93], is a level surface of N, which in turn means that $^0_1$ and $^0_2$ both bound product neighborhoods of degenerate ends of M on both sides, one contained in D and the other in E. Hence, M is a doubly degenerate hyperbolic metric on S R.

For the second part of the proposition, we need to show that a Borel subset of H^3 is dened if we require (M; p) to satisfy the following three conditions:

(A) M has no cusps,
(B) $_1$M is nitely generated,
(C) p lies in an R-core of M.

We'll rst show that (A) and (B) each dene Borel subsets, and deal with (C) afterwards. Note that (C) doesn't make sense on its own, since R-cores are only dened for manifolds with nitely generated $_1$ and no cusps.

To see that (A) denes a Borel set, we check that for each R > 0 and small > 0, the set of all (M; p) where there is a cuspidal -thin part at distance at most R from p is closed, and then take a union over R 2 N. For if (M_i; p_i) ! (M_1; p_1), we can write M_i = H^3= $_i$ in such a way that $_i$! $_1$ in the Chabauty topology on subgroups of PSL_2C, and the points p_i are all projections of a xed p̃ 2 H^3, see [76, Chapter 7]. Cuspidal -thin parts at distance at most R from p_i 2 M_i then give elements 1 = $_i$ 2 $_i$ such that

1) jtr $_i$j = 2 for all i,
2) there are points $\tilde{x}_i$ 2 H^3 with d($\tilde{x}_i$; p̃) R and d($\tilde{x}_i$; $_i$($\tilde{x}_i$)) 2.

After passing to a subsequence, $\tilde{x}_i$! $\tilde{x}_1$ 2 H^3 and $_i$! $_1$ 2 $_1$, where d($\tilde{x}_1$; p̃) R and jtr $_1$j = 2. Passing to the quotient, we have a cuspidal -thin part at distance at most R from p_1 2 M_1.

To prove that (B) is a Borel condition, we use:

Lemma 6.3. Fix a compact 3-manifold N_0, a Riemannian metric on N_0, and a constant > 1. Let S be the set of all (M; p) 2 H^3 that admit a smooth embedding f : N_0 ,! M such that

1) the iterated derivative maps D^kf : T^kN_0 ! T^kM are locally -bilipschitz on the 1-neighborhood of the zero section in T^kN_0, for k = 0; 1; 2,
2) the image N = f(N_0) contains p,
3) each component of M r N is homeomorphic to S R, for some closed surface S.

Then S is a Borel subset of H^3.

See §A.2 for details about iterated derivatives. The point of 1) is the following, though. As shown in at the end of the proof of Theorem A.10, bounds on iterated total derivatives up to order 2 give bounds in fixed local coordinates for the C^2-norm of f. So by Arzela-Ascoli, if $f_i : N_0 \hookrightarrow M$ is a sequence of embeddings satisfying 1) and 2), then after passing to a subsequence we may assume that f_i converges in the C^1-topology to some C^1-embedding $f : N_0 \to M$. Working in normal-bundle coordinates in a regular neighborhood of $\partial f(N_0)$, we can then move any f_i (with i large) by a small C^1-isotopy so that the image agrees with $f(N_0)$. In particular, the images of all such f_i differ by small isotopies.

Proof of Lemma 6.3. We'll write S for the set of all (M, p) admitting an embedding as above. Fix $R > 0$ and consider the set S_R of all (M, p) such that there is an embedding $N_0 \hookrightarrow M$ satisfying 1) and 2), and also a compact submanifold $N' \subset M$ that contains the radius R-ball around p, and where every component of $N' \setminus \mathrm{int}(N)$ is homeomorphic to $S \times [0, 1]$ for some closed surface S. This S_R is open in the smooth topology, since the approximate isometries defining the smooth topology (see §A.1) allow us to transfer compact submanifolds of (M, p) to nearby (N, q), with small metric distortion.

We claim that $S = \cap_R S_R$, which will imply S is Borel. The forward inclusion is obvious, so assume (M, p) is in the intersection. Then there is a sequence $R_i \to \infty$ such that if $f_i : N_0 \hookrightarrow M$, N_i and N_i^0 are the corresponding embeddings and submanifolds, we have $N_1^0 \subset N_2^0 \subset \dots$. Passing to a subsequence, we may assume by 1) and 2) that the embeddings $f_i : N_0 \hookrightarrow M$ all differ by small isotopies. Hence, the complement of $\mathrm{int}(N_1)$, say, in every N_i^0 is a union of topological products $S \times [0, 1]$. Consequently, each $N_{i+1}^0 \setminus \mathrm{int}(N_i^0)$ is also a union of products $S \times [0, 1]$, so taking a union over i, the components of $M \setminus N_1$ are homeomorphic to $S \times \mathbb{R}$.

To prove (B) defines a Borel set, apply the lemma and take a union over the countably many homeomorphism types of compact 3-manifolds N_0, over a countable dense subset of the space of Riemannian metrics on a given N_0, and over $\epsilon \in \mathbb{N}$. This proves that a Borel subset is defined by the condition that there is a submanifold $N \subset M$ containing p whose complementary components are homeomorphic to products $S \times \mathbb{R}$. By the Tameness Theorem, this is equivalent to $\pi_1 M$ being finitely generated.

The proof that (C) defines a Borel set uses a more complicated version of Lemma 6.3, but the argument is very similar. Fixing a compact 3-manifold N_0, a Riemannian metric on N_0, and a constant $\epsilon > 1$, let S be the set of all $(M, p) \in \mathcal{H}^3$ with no cusps and finitely generated π_1, that admit a smooth embedding $N_0 \hookrightarrow M$ satisfying 1) and 2) of Lemma 6.3, and whose image is an R-core of M.

We claim that S is Borel. To that end, fix $T > 0$ and consider the set S_T of all (M, p) with finitely generated π_1 and no cusps, satisfying the following conditions. We require that there is an embedding $N_0 \hookrightarrow M$ with bilipschitz constant less than ϵ, whose image N contains p and satisfies

1) the diameter of N is less than R,
2) N is contained in an open radius 1 neighborhood of the convex core of M.

We also require that N is contained in a compact submanifold $N' \subset M$ whose interior contains the closed radius T ball around N, such that

3) each component E of $N' \setminus \mathrm{int}(N)$ is homeomorphic to $S \times [0, 1]$, for some closed surface S,

4) if a component E of N^0 n int(N) intersects the convex core of M, then through each point p 2 E that lies more than an -electric distance of R from M n E, there is a level surface in E = S [0; 1] that passes through p that has -electric diameter less than R.

We now claim that S_T is Borel. As we have shown above, the conditions that M has no cusps, and that ${}_1$M is nitely generated are both Borel. Furthermore, the approximate isometries dening the smooth topology (see §A.1) allow us to transfer compact submanifolds of (M; p) to nearby (N; q), with small metric distortion. So, the existence of an embedding N_0 ,! M as above that satises 1) and 3) is an open condition.

To deal with the other conditions, let N > 0 and let H^3_N be the set of all (M; p) 2 H^3 where the injectivity radius of (M; p) is bounded above by N throughout the convex core of M. Each H^3_N is a closed subset of H^3, and the convex core of M varies continuously as (M; p) varies within H^3_N, by [77, Proposition 2.4]. Using this and the previous paragraph, one can see that 2) and 4) are also open conditions within H^3_N, so S_T \ H^3_N is a Borel subset of H^3_N for all N. As any hyperbolic 3-manifold with nitely generated ${}_1$ has an upper bound for the injectivity radius over its convex core, [32, Corollary A], this expresses S_T as a union of Borel sets, so S_T is Borel.

We claim that S = $\backslash_T$ S_T, which will imply S is Borel. The forward inclusion is obvious, so assume (M; p) is in the intersection. Arguing as in the proof of Lemma 6.3, we may assume that we have a xed embedding f : N_0 ,! M satisfying 1) and 2), and a sequence T_i ! 1 such that there are N^0_i whose interiors contain the closed T_i-ball around N = f(N_0) and satisfy 3) and 4). Again as in the proof of Lemma 6.3, the complementary components of N are all homeomorphic to products S R. But then conditions 3) and 4) in the denition of an R-core follow for N from condition 4) above. For if a component of M r N intersects the convex core, it must have the necessary level surfaces with bounded electric diameter (and is a neighborhood of a degenerate end). Otherwise, we have condition 4) in the denition of an R-core.

This proves that S is Borel, which proves (C) is a Borel condition, and Proposition 6.2 follows.

6.0.2. The case with cusps. In the presence of cusps, the proof is the same, but more complicated. When M has nitely generated ${}_1$, the topological ends of M may be split by cusps into 'geometric ends'. More precisely, xing > 0 we let M_{np} be the manifold with boundary that is the complement of the cuspidal -thin part of M. Then the topological ends of M_{np} are called geometrically nite or degenerate depending on whether they have neighborhoods disjoint from, or contained in, the convex core of M. The denition of an R-core is basically the same, except that M should be replaced by M_{np}, complementary components should now be homeomorphic to S R, where S is a surface with boundary, and the dening electric distance should be chosen smaller than that dening M_{np}.

Proposition 6.2 should now allow cusps, and exclude in 1) products S R, where S is a nite type surface. Its proof is the same, as long as one keeps track of the relationship between the new denition of R-cores and M_{np}. However, the increase in the already formidable amount of notation will be unpalatable.

Appendix A. Spaces of Riemannian manifolds

In this appendix, we discuss the smooth topology on the space of pointed Riemannian manifolds, and related topologies on similar spaces.

A.1. Smooth convergence and compactness theorems. As introduced in the introduction, let $\mathcal{M}^d$ be the set of isometry classes of connected, complete, pointed Riemannian manifolds (M, p).

Remark A.1. The class of all pointed manifolds is not a set. However, every connected manifold has at most the cardinality of the continuum, so one can discuss $\mathcal{M}^d$ rigorously by only considering manifolds whose underlying sets are identied with subsets of $\mathbb{R}$. The same trick works in all the related spaces below, so we will make no more mention of set theoretic technicalities.

A sequence of pointed Riemannian manifolds (M_i, p_i) C^k-converges to (M, p) if there is a sequence of C^k-embeddings

$$f_i : B(p, R_i) \longrightarrow M_i \tag{38}$$

with $R_i \to \infty$ and $f_i(p) = p_i$, such that $f_i^* g_i \to g$ in the C^k-topology, where g_i, g are the Riemannian metrics on M_i, M. Here, C^k-convergence of tensors is dened locally: in each pre-compact coordinate patch, the coordinates of the tensors and all their derivatives up to order k should converge uniformly. Note that each metric $f_i^* g_i$ is only partially dened on M, but their domains of denition exhaust M, so it still makes sense to say that $f_i^* g_i \to g$ on all of M. We say that $(M_i, p_i) \to (M, p)$ smoothly if the convergence is C^k for all $k \in \mathbb{N}$.

Whether the convergence is C^k or smooth, we will call such an (f_i) a sequence of almost isometric maps witnessing the convergence $M_i \to M$. When the particular radii R_i do not matter, we will denote our partially dened maps by

$$f_i : M \dashrightarrow M_i,$$

where the notation indicates that each f_i is only partially dened on M, but any point $p \in M$ is in the domain of f_i for all large i.

Of course, another way to dene C^k-convergence would be to require that for every xed radius $R > 0$, there is a sequence of maps $f_i : B(p, R) \longrightarrow M_i$ satisfying the properties above. To translate between the two denitions, we can restrict the f_i in (38) from R_i-balls to R-balls, or in the other direction, we can take a diagonal sequence where R increases with i. Most of the time, we will use the 'xed R' perspective in this appendix.

In some references, e.g. Petersen [81, §10.3.2], R is xed and the maps f_i are dened on open sets containing $B(p, R)$ and their images are required to contain $B(p_i, R) \subset M_i$. Such restrictions on the image of the f_i do not change the resulting C^k-topology, though, since for large i the maps f_i are locally 2-bilipschitz embeddings, and we can appeal to the following Lemma.

Lemma A.2. *Suppose M, N are complete Riemannian d-manifolds, $p \in M$. If $f : B(p, R) \longrightarrow N$ is a smooth locally λ-bilipschitz embedding,*

$$f(B(p, R)) \supset B(f(p), R/\lambda).$$

Proof. Fix a point $q \in B(f(p), R/\lambda)$ and let $\gamma : [0, 1] \longrightarrow N$ be a length-minimizing geodesic with $\gamma(0) = f(p)$ and $\gamma(1) = q$. Let

$$T = \{t \mid \exists \, \tilde\gamma_t : [0, t] \longrightarrow B(p, R) \text{ with } \tilde\gamma_t(0) = p,\ f \circ \tilde\gamma_t = \gamma|_{[0,t]}\} \subset [0, 1].$$

The set T is a subinterval of $[0;1]$ that contains 0. It is open in $[0;1]$, since f is a local dieomorphism. We claim that it is also closed. First, as f is a local dieomorphism the lifts $_t$ are unique if they exist, so if $t\ 2\ [0;1]$ is a limit point of T then the limiting lifts patch together to give

$$: [0;t) \quad ! \quad R; \quad (0) = p; \ f \quad = j_{[0;t)}:$$

The path is itself locally -bilipschitz, so its image is contained within the compact subset $\overline{B(p; d(f(p); q))}\ B(p;R)$. From this and the fact that it is locally bilips-chitz, can be continuously extended to a map $[0;t] \quad ! \quad R$. This map must lift , so $t\ 2\ T$. Therefore, $T = [0;1]$, implying in particular that $q = (1)\ 2\ f(B(p;R))$.

One reason that these convergence notions are useful is that sequences of manifolds with 'uniformly bounded geometry' have convergent subsequences.

Denition A.3. Suppose that M is a complete C^k-Riemannian manifold. A subset A M has C^k-bounded geometry if for some xed $r; ; L > 0$ there is a system of coordinate charts

$$_s : B_{R^d}(0;r) \quad ! \quad U_s \quad M$$

with the following properties:

1) for every $p\ 2\ A$, the ball $B(p;)\ \ U_s$ for some s,
2) $_s$ is locally L-bilipschitz,
3) all coordinates of the metric tensor $(g_s)_{ij}$ have C^k-norm at most L,
4) the transition maps $^1\ _t\ _s$ have C^{k+1}-norm at most L.

A sequence of subsets $A_i \quad M_i$ of complete Riemannian manifolds has uniformly C^k-bounded geometry if the constants $r; ; L$ can be chosen independently of i.

The conditions above are a simplication of those given by Petersen [81, pp. 289, 297] that are sucient for the following theorem.

Theorem A.4. Suppose that $(M_i; p_i)$ is a sequence of complete pointed C^k-Riemannian manifolds and that for some xed $R > 0$ the balls $B(p_i;R)\ \ M_i$ have uniformly C^k-bounded geometry. Then there is a complete pointed $C^{k\ \ 1}$-Riemannian manifold $(M;p)$ and, for suciently large i, embeddings

$$f_i : B(p; R \quad 1) \quad ! \quad M_i$$

with $f_i(p) = p_i$ such that $f_i g_i \ ! \ g$ in the $C^{k\ \ 1}$-topology on $B(p; R \quad 1)$, where $g_i; g$ are the Riemannian metrics on $M_i; M$.

In other words, if the balls $B(p_i;R)\ 2\ M_i$ have uniformly C^k-bounded geometry, there is a subsequence on which we see $C^{k\ \ 1}$-convergence of the $(M_i; p_i)$, at least within a distance of $R \quad 1$ from the base points. This is a version of Cheeger's compactness theorem [81, Ch. 10, Thm. 3.3] for R-balls { Cheeger's theorem usually requires uniform geometry bounds over the entire manifolds and then gives a fully convergent subsequence.

Proof of Theorem A.4. The theorem is not implied by the statement of Cheeger's compactness theorem given in [81, Ch. 10, Thm. 3.3], since the latter requires uniform C^k bounds over the entire manifolds $(M_i; p_i)$, but the proof is the same.

The ideas to take a subsequential Gromov-Hausdor limit $(X; x)$ of the balls $B_{M_i}(p_i;R)$, which exists by the uniform geometry bounds. An atlas of $C^{k\ \ 1}$-charts for X is obtained as a limit of the charts for $B_{M_i}(p_i;R)$ with uniformly C^k-bounded geometry, and then one shows that the convergence is $C^{k \quad 1}$ in addition to Gromov-Hausdor. Then choose a complete pointed Riemannian manifold $(M; p)$ such that $B(p; R \quad 1) \quad M$ and $B(x; R \quad 1) \quad X$ are isometric.

The details are entirely the same as those of [81, Ch. 10, Thm. 3.3].

Theorem A.4 also gives a strong version of Cheeger's theorem in which the uniform geometry bounds may depend on the distance to the base point:

Corollary A.5. Let $(M_i; p_i)$ be a sequence of complete pointed Riemannian manifolds such that for every $R > 0$ and k 2 N, the balls $B(p_i; R) \quad M_i$ have uniformly C^k-bounded geometry, where the bounds may depend on $R; k$ but not on i. Then $(M_i; p_i)$ has a smoothly convergent subsequence.

Proof. Applying a diagonal argument and passing to a subsequence, we may assume that for every k 2 N, there is a complete pointed $C^{k \quad 1}$-Riemannian manifold $(L_k; q_k)$ and, for suciently large i, embeddings

$$f_{i;k} : B(p; k \quad 1) \quad ! \quad M_i$$

with $f_i(q_k) = p_i$ such that $f g_i \; ! \; g_k$ in the $C^{k \quad 1}$-topology on $B(p; R \quad 1)$, where $g_i; g_k$ are the Riemannian metrics on $M_i; M$.

By Arzela-Ascoli's theorem, for each k there is a pointed isometry

$$B_{L_k}(q_k; k \quad 1) \quad ! \quad B_{L_{k+1}}(q_{k+1}; k \quad 1):$$

So, the direct limit of the system

$$B_{L_2}(q_2; 1) \quad ! \quad B_{L_3}(q_3; 2) \quad ! \quad B_{L_4}(q_4; 3) \quad !$$

is a complete pointed C^1-Riemannian manifold to which $(M_i; p_i)$ smoothly converges.

A.2. Metrizability of M^d in the smooth topology. The goal here is to show that smooth convergence on M^d is induced by a Polish topology. As mentioned in the introduction, this result was established independently and concurrently by Candel, Alvarez Lopez and Barral Lijo [10]. Their paper became available earlier than ours, so the theorem is certainly theirs. The two approaches use the same key idea, encoding partial derivatives of a metric in iterated Sasaki metrics, but ours produces an explicit metric that will be used elsewhere in the paper, namely in §A.5. Our approach is also a bit simpler, since we use metric neighborhoods of the zero section in iterated tangent bundles instead of the iteratively dened neighborhoods in [10].

Suppose M is a manifold with a Riemannian metric g. Sasaki [86] introduced a Riemannian metric g^1 on the tangent bundle TM. If $(x_1; : : : ; x_d)$ is a coordinate system for some $U \quad M$, let

$$(x; v) = (x_1; : : : ; x_d; v_1; : : : ; v_d)$$

be the induced coordinates on TU, where $v_i = Dx_i$. At a point $(x; v)$ 2 TU, the Sasaki metric g^1 is given as follows, see [86, p. 342]: for $1 \quad i; j \quad d$,

$$(g^1)_{ij} = g_{ij} + \quad X \quad g \;\; {}_i \;\; {}_j vv \, {}_{1;;;d}$$

(39) $$(g^1)_{i(d+j)} = \quad X \quad {}_{ji} v \, {}_{1d}$$

$$(g^1)_{(d+i)(d+j)} = g_{ij}.$$

Here, $_{kaj}$ and ^{aj} are the Christoel symbols of g of the rst and second kind, and all the metric data on the right sides of the equations are taken at $x \in U$.

The k-fold iterated tangent bundle of a manifold M is the manifold

$$T^kM = T(T(T(M))).$$

Any smooth map f of manifolds induces a smooth map of iterated tangent bundles, the iterated total derivative D^kf, and if M has a Riemannian metric g then T^kM inherits the k-fold iterated Sasaki metric g^k. Note that

$$T^kM \to T^{k-1}M \to \cdots \to TM \to M$$

is a sequence of vector bundles, so for each i we have a zero section

$$: T^{i-1}M \to T^iM.$$

Dene the zero section of the ber bundle $T^kM \to M$ to be the iteration ^{k}

$$: M \to T^kM.$$

We now show that the iterated Sasaki metric g^k encodes all order $\leq k$ derivatives of g. We will state this in a strong way that will be useful in Corollary A.7, but the point is that partial derivatives of the g_{ij} can be written in terms of the g^k.

Lemma A.6. Suppose g is a Riemannian metric on an open subset $U \subset \mathbb{R}^d$ and let g^1 be the Sasaki metric on TU. Fix coordinates $(x_1, \dots, x_d)$ on U and let $(x_1, \dots, x_d, v_1, \dots, v_d)$ be the induced coordinates on TU, where $v_i = Dx_i$.

For every $x = (x_1, \dots, x_d) \in U$ and indices $, i, j \in \{1, \dots, d\}$, we have

(40) $$g_{ij}(x) = g^1_{(i+d)(j+d)}(x, v), \quad \forall v \in T_xU$$

(41) $$\partial g_{ij}(x) = t\,(g^1_{i(d+j)} + g^1_{j(d+i)})(x, 0, \dots, \underbrace{\tfrac{1}{t}}_{\text{th place}}, \dots, 0), \qquad \forall t > 0.$$

Now fix k, and iterate the above construction to give a system of coordinates on T^kU. For any compact $C \subset U$ and any open neighborhood $O \supset {}^k(C)$ in T^kU, any partial derivative $\partial_{1,\dots,l}\, g_{ij}(x)$ with $x \in C$, $0 \leq l \leq k$ can be represented as a linear combination

(42) $$\partial_{1,\dots,l}\, g_{ij}(x) = \sum t_n\, g^k_{nn}(v_n)\, n$$

where $v_n \in O$ and $1 \leq n, n \leq kd$. Here, v_n, n, n, t_n are determined just by the indices $_1, \dots, _l, i, j$ and the choices of C, U, and not by the metric g.

Proof. The $k = 1$ case follows from (39), using the identity $_{ji} + _{ij} = \partial g_{ij}$. For (42), rst choose some $t > 1$ such that for all $x \in C$, the point

$$(x, 0, \dots, \underbrace{\tfrac{1}{t}}_{\text{th place}}, \dots, 0) \in O.$$

Then (42) is proved inductively: every time a successive partial derivative of is taken, one can instead consider the appropriate linear combination of entries of the next Sasaki metric, given in (40). One takes at most k of these derivatives, and then uses (41) to encode the resulting data in the k^{th} iterated Sasaki metric, instead of a previous one.

It will follow that C^k-convergence of metrics g_i is equivalent to C^0-convergence of the Sasaki metrics g_i^k. Here, recall that C^k-convergence means uniform convergence on compact sets of all derivatives up to order k. In fact, slightly more is true.

Corollary A.7. Fix $k \in \mathbb{N}$. Suppose that g_i, g are Riemannian metrics on some manifold M, and let g_i^k, g^k be the induced Sasaki metrics on T^kM. Then the following are equivalent:

1) $g_i \to g$ in the C^k-topology on M,
2) $g_i^k \to g^k$ in the C^0-topology on T^kM,
3) for any compact $C \subset M$, there is some open set $O \subset T^kM$ containing $\pi^k(C)$ such that $g^k_i \to g^k$ uniformly on K.

Proof. Covering M with a locally finite set of coordinate charts, it suffices to prove the Corollary when M is an open subset of $\mathbb{R}^n$. So, we will feel free to use the coordinate expressions in (39) and Lemma A.6 below without comment.

For (1) implies (2), fix a compact subset $K \subset T^kM$. The C^k-convergence $g_i \to g$ implies that all derivatives up to order k of g_i converge to those of g, uniformly on the projection of K into M. In particular, when computing each successive Sasaki metric, the Christoffel symbols in (39) converge. Since K is compact, all coordinates v, ν from (39) that are relevant in the construction of g^k are bounded. It follows that $g^k_i \to g^k$ uniformly on K.

Since C^0-convergence is by definition uniform convergence on compact sets, the implication (2) $\Rightarrow$ (3) is obvious, taking O to be any such open subset that has compact closure in T^kM.

It remains to prove (3) $\Rightarrow$ (1): Given a compact subset $C \subset M$, let O be as in (3). For each $x \in M$, we can use (42) to write the partial derivatives of each g_i as linear combinations of entries of the g^k_i, evaluated at points of O, where then particular linear combination is independent of i, and also works for g. These linear combinations converge uniformly, by our assumption, so $g_i \to g$ in the C^k-topology as desired.

The advantage of using convergence of Sasaki metrics instead of C^k convergence is that it is easier to metrize, since two Sasaki metrics can then be compared by analyzing the minimal distortion of a bilipschitz map between them. Now, it is not the case that when two metrics on M are C^k-close, the associated metrics on T^kM are bilipschitz. (We thank a referee for first bringing this to our attention.) For instance, looking at (39), if two metrics are C^1-close then their Christoffel symbols are close, but the expression given for the Sasaki metric also includes the coordinates v, ν, which can be arbitrarily large[13]. To fix this, we only look at the bilipschitz distortion on subsets of T^kM in which the coordinates v, ν are bounded. Namely, for each $r > 0$ and $k \in \mathbb{N}$, let

$$Z^k_r(M) \subset T^kM$$

[13] One can also see this issue when calculating $D^2 f$, where $f : \mathbb{R} \to \mathbb{R}$ is a map between open subsets of $\mathbb{R}$ that is C^2-close to the identity. On $T^2\mathbb{R}$, the map $D^2 f$ has the form

$$(x, v, w) \mapsto \left(x, f'(x)v, \begin{pmatrix} f'(x) & 0 \\ f''(x)v & f'(x) \end{pmatrix} w\right),$$

where here $v \in T\mathbb{R}_x = \mathbb{R}$ and $w \in T(T\mathbb{R})_{(x,v)} = \mathbb{R}^2$. The term $f''(x)v$ can be arbitrarily large, so $D^2 f$ is only bilipschitz if $f''(x) \equiv 0$.

be the radius-r neighborhood of the zero section $\sigma^k(M)$, with respect to the k-fold iterated Sasaki metric. Then we have:

Corollary A.8. Suppose that $(M_n, p_n) \in \mathcal{M}^d$, where $n = 1, 2, \ldots$. Then (M_n, p_n) converges C^k to $(M, p) \in \mathcal{M}^d$ if and only if for every $R > 0$, we have that for sufficiently large n, there are smooth embeddings

$$f_n : B_M(p, R) \to M_n$$

with $f_n(p) = p_n$ such that for some $r > 0$, the iterated total derivatives

$$D^k f_n : T^k B_M(p, R) \to T^k M_n$$

are locally λ_n-bilipschitz embeddings on $Z_r := Z_r(\sigma^k B_M(p, R))$, with $\lambda_n \to 1$.

Proof. Let g and g_n be the metrics on M and M_n. The point is to show that the condition that $D^k f_n$ is locally λ_n-bilipschitz on Z_r, with $\lambda_n \to 1$, is equivalent to the condition that $f_n^* g_n \to_n g$ in the C^k-topology. First, note that

$$(D^k f_n)^* g_n^k = (f_n^* g_n)^k,$$

where g_n^k is the Sasaki metric on $T^k M_n$ and $(f_n^* g_n)^k$ is the Sasaki metric on $T^k B_M(p, R)$ corresponding to $f_n^* g_n$. So, $D^k f_n$ is locally λ_n-bilipschitz on Z_r exactly when

$$\mathrm{id} : (B_M(p, R), g^k) \to (B_M(p, R), (f_n^* g_n)^k) \tag{43}$$

is locally λ_n-bilipschitz on Z_r.

The 'only if' direction then follows from our previous work, since if the map (43) is locally λ_n-bilipschitz on $Z_r(\sigma^k M)$ with $\lambda_n \to 1$, then

$$(f_n^* g_n)^k \to g^k$$

uniformly on Z_r, implying that $f_n^* g_n$ converges C^k to g by Corollary A.7.

The 'if' direction is also easy. Suppose that $f_n^* g_n \to g$ in the C^k-topology. Then Corollary A.7 implies that $(f_n^* g_n)^k \to g^k$ in the C^0 topology. The ball $B_M(p, R)$ has compact closure in M, so the set Z_r has compact closure in $T^k M$. Hence, $(f_n^* g_n)^k \to_n g^k$ uniformly on Z_r, implying that the identity map in (43) is locally λ_n-bilipschitz with $\lambda_n \to 1$. $\square$

Finally, we record for later use the following fact.

Fact A.9. Fix $\lambda > 1$, $k \in \mathbb{N}$ and $r > 0$. Suppose that $f : M \to N$ is a smooth map of Riemannian manifolds and that $D^k f$ is locally λ-bilipschitz when restricted to some neighborhood of the zero section $\sigma^k(M) \subset T^k M$, for instance the neighborhoods $Z_r(\sigma^k M)$ above. Then the map f is itself locally λ-bilipschitz.

Proof. Let $x \in M$, set $v_0 := x$ and set $v^i = \sigma^i(x)$, for $i = 1, \ldots, k$. For each i, we have isometric embeddings

$$\iota : T^i M_{v_{i-1}} \cong T((T^i M)_{v_{i-1}})_{v_i} \hookrightarrow T^{i+1} M_{v_i},$$

which satisfy the commutative diagram

$$\begin{array}{ccc} T^i M_{v_{i-1}} & \longrightarrow & T^{i+1} M_{v_i} \\ \downarrow D^i f & & \downarrow D^{i+1} f \\ T^i N_{D^{i-1} f(v_{i-1})} & \longrightarrow & T^{i+1} N_{D^i f(v_i)} \end{array};$$

dening similarly on N. The commutativity is just the statement that a linear map is its own derivative, and the fact that is an isometric embedding is immediate from the denition of the Sasaki metric. Composing, we have a diagram

$$\begin{array}{ccc} TM_x & \xrightarrow{\ k\ } & T^{k+1}M_{v_k} \\ \downarrow{\scriptstyle Df} & & \downarrow{\scriptstyle D^{k+1}f} \\ TN_{f(x)} & \xrightarrow{\ k\ } & T^{k+1}N_{D^kf(v^k)} \end{array} ;$$

so if the linear map $D^{k+1}f$ is -bilipschitz, so must be Df.

Corollary A.8 suggests a description of a basis of neighborhoods around a point $(M; p)$ 2 M^d. We dene the k^{th}-order $(R; r;)$-neighborhood of $(M; p)$, written

$$N^k_{R;r;}(M; p);$$

to be the set of all $(N; q)$ such that there is a smooth embedding

$$f : B_M(p; R) \; ! \; N$$

with $f(p) = q$ such that $D^kf : T^kU \; ! \; T^kN$ is locally -bilipschitz with respect to the iterated Sasaki metrics on $Z^k_r(B_M(p; R))$. Note that $N^k_{R;r;}(M; p)$ is a closed neighborhood of $(M; p)$. By Arzela-Ascoli, its interior is the open neighborhood

$$N^k_{R;r;}(M; p)$$

that is dened similarly, except that we require D^kf to be locally ^{0}-bilipschitz for some $^0 <$.

Theorem A.10. M^d has the structure of a Polish space (a complete, separable metric space), in which convergence is smooth convergence.

Proof. For each $R > 0$ and k 2 $\mathbb{N}$, dene a function $d_{R;k} : M^d \; M^d \; ! \; \mathbb{R}$ by

$$d_{R;k}((M; p); (N; q)) = \inf f\log \; j \; (N; q) \; 2 \; N^k_{R=;1=;}(M; p)g:$$

Each $d_{R;k}$ satises an (asymmetric) triangle inequality. For suppose we have $(M; p); (N; q); (Z; z)$ 2 M^d and basepoint respecting embeddings

$$f : B_M(p; R=) \; ! \; N; \quad g : B_N(q; R=) \; ! \; Z$$

such that D^kf and D^kg are locally -bilipschitz and locally -bilipschitz embeddings, respectively, when restricted to the sets

$$Z^k_{1=}(B_M(p; R=)) \;\; T^kM; \quad Z_{1=}(B^k_N(q; R=)) \;\; T^kN:$$ By Fact A.9, the map f is also a locally -bilipschitz embedding, so

$$f \quad B_M \quad p; \frac{R}{} \qquad B_N \quad q; \frac{R}{} \quad ;$$

and since D^kf is -bilipschitz,

$$D^kf \quad Z^k_{1=()}(B_M(p; R=())) \quad Z_{1=}(B_N(^k_q; R=)):$$

Therefore, the composition $gf : B_M(p;) \xrightarrow{R} N$ is dened and the map $D^k(gf)$ is locally $()$-bilipschitz on $Z_{1=()}(B^k_M(p; R=()):$ So,

$$d_{R;k}((M; p); (Z; z)) \;\; \inf_{;} \log$$

$$= \inf \log + \inf \log$$

$$= d_{R;k}((M;p);(N;q)) + d_{R;k}((N;q);(Z;z)):$$

The subsets of M^d dened for each $(M;p) \in M^d$, $R > 0$, $k \in N$, > 0 by

$$d_{R;k}((M;p);) <$$

form a basis for a smooth topology on M^d that induces smooth convergence, by Lemma A.8. Although the $d_{R;k}$ are not symmetric, the reversed inequalities

$$d_{R;k}(;(M;p)) <$$

dene a basis for the same topology, as Lemma A.2 allows the relevant locally bilipschitz maps to be inverted at the expense of decreasing R. So, the smooth topology is generated by the family of pseudo-metrics

$$\hat{d}_{R;k} : M^d \times M^d \to R; \quad \hat{d}_{R;k}(x;y) = d_{R;k}(x;y) + d_{R;k}(y;x):$$

As the topology on M^d induced by a particular pseudo-metric $\hat{d}_{R;k}$ becomes ner if $R;k$ are increased, it suces to consider only $\hat{d}_{k;k}$ for $k \in N$. Therefore, the following is a metric on M^d that induces the smooth topology:

$$D : M^d \times M^d \to R; \quad D\ x;y = \sum_{k=1}^{\infty} 2^{-k} \min\{\hat{d}_{k;k}(x;y);1\}:$$

We now show that M^d is separable. An element $(M;p) \in M^d$ is a limit of closed Riemannian manifolds: for instance, we can exhaust M by a sequence of compact submanifolds with boundary, double each of these and extend the Riemannian metric on one side arbitrarily to the other. So, it suces to construct a countable subset of M^d that accumulates onto every closed manifold in M^d.

There are only countably many dieomorphism types of closed manifolds: this is a consequence of Cheeger's niteness theorem [37], for instance. So, it suces to show that the space $M(M)$ of isometry classes of pointed closed Riemannian manifolds in the dieomorphism class of some xed M is separable in the smooth topology. This space $M(M)$ is a continuous image of the product of M with the space of Riemannian metrics on M, with the smooth topology on tensors. The manifold M is separable, and so is the space of Riemannian metrics on M, by the Weierstrass approximation theorem. So, their product is separable, implying $M(M)$ is too, nishing the proof.

Finally, we want to show that $(M^d;D)$ is complete, so let $(M_i;p_i)$ be a Cauchy sequence. We claim that for every $R > 0$ and $k \in N$, the balls

$$B(p_i;R) \subset M_i$$

have uniformly C^k-bounded geometry, in the sense of Denition A.3. Corollary A.5 will then imply that $(M_i;p_i)$ has a smoothly convergent subsequence, which will nish the proof.

It suces to show that there are arbitrarily large $R;k$ for which the balls

$$B(p_i;R) \subset M_i$$

have uniformly C^k-bounded geometry. So, x some $k \in N$. Since $(M_i;p_i)$ is D-Cauchy and $D = \sum_{j=1}^{\infty} 2^{-j} \min\{\hat{d}_{j;j};1\}$; the $d_{k;k}$-diameter of the tail of $(M_i;p_i)$

can be made arbitrarily small. In other words, there is some $(M;p) \in \mathcal{M}^d$ such that for suciently large i,

$$d_{k;k}\ (M;p);(M_i;p_i)\ <\ \log 2:$$

This means that for each large i there is a pointed, smooth embedding

$$f_i : B_M(p;k) \to M_i$$

such that $D^k f_i$ is locally 2-bilipschitz on $Z := Z^k_{\frac{1}{2}}(B_M(p;k=2))$. By Lemma A.2,

$$f(B_M(p;k=2))\ B_{M_i}(p_i;k=4):$$

By precompactness, the ball $B_M(p;k=2\ \ 1)\ M$ has C^k-bounded geometry, in the sense of Denition A.3, for some constants $r;;L$. We will use the maps f_i to translate this to uniform C^k-geometry bounds for the balls

$$B_{M_i}(p_i;R)\ M_i;\ \text{where } R := (k\ \ 1)=4:$$

Let $_s : B \to U_s\ M$ be nitely many coordinate charts as in Denition A.3, where B is a ball around the origin in $\mathbb{R}^d$, and by shrinking B assume that $U_s\ B_M(p;k)$ for all s. Dene

$$_{i;s} : B \to U_{i;s}\ M_i;\ \ _{i;s} = f_i\ _s\ \ \forall i;s:$$

It is now straightforward to verify that 1) – 4) of Denition A.3 are satised for the subsets $B_{M_i}(p_i;R)\ M_i$ by the charts $_{i;s}$, with modied constants. As the maps $f_i : B_M(p;R) \to M_i$ are locally 2-bilipschitz, the $_{i;s}$ are locally 2L-bilipschitz, so by Lemma A.2 the $=2$-ball around every $q \in B_M(p;R^0)$ is contained in some $U_{i;s}$, as $R^0 = (R\ \ 1)=2$. The transition maps of 4) are unchanged by the composition, so it remains to prove condition 3). The argument for here is like an eective version of (3) $\Rightarrow$ (1) in Corollary A.7. Fix some s, and let $O = D^k\ {}_s^{-1}(Z)\ T^kB$, where $Z\ \ T^kM$ is the precompact subset dened above on which f_i is locally 2-bilipschitz. By Lemma A.6, each partial derivative of a Riemannian metric g on B can be expressed as some xed linear combination

$$\partial_{1;\dots;l}\, g_{ij}(x) = \sum t_n\ g^k_{n\,n}(v_n)\ _n \tag{44}$$

of coecients of the metric g^k on T^kB, where $v_n \in O$ and the linear combination depends only on the partial derivative taken and the neighborhood O, not the metric g. So, let g_M and g_i be the metrics on M and M_i, respectively. By assumption, the coecients of the metric g^M_s have bounded C^k norm, so since $O\ \ T^kB$ is precompact, the construction of the Sasaki metric implies that all the coecients of $(g_M\ _s)^k = (D^k{}_s)g^k_M$ are bounded on O. But since D^kf_i is 2-bilipschitz on Z, the metrics g^k_M and $(D^kf_i)g^k_i$ dier by a factor of at most 2 on Z, and hence the coecients of the metric

$$((f_i\ _s)g_i)^k = (D^k{}_s)((D^kf_i)g^k_i)$$

are also bounded on O. Equation (44) then implies that all partial derivatives of the metric $(f_i\ _s)g_i$ are bounded, which proves condition (3).

This shows that for all $k \in \mathbb{N}$, the balls $B(p_i;(k\ \ 1)=2)\ M_i$ have uniformly C^k-bounded geometry. (We showed this explicitly for large i, but the initial nitely many terms only contribute a bounded increase to the constants.) So, Corollary A.5 implies that $(M_i;p_i)$ has a smoothly convergent subsequence.

A.3. Vectored and framed manifolds. It is often convenient to supplement the basepoint of a pointed manifold (M, p) with additional local data: for instance, a unit vector or an orthonormal basis for T_pM. An orthonormal basis for T_pM is called a frame for M at p, and we let $\mathcal{F}M \to M$ be the bundle of all frames for M.

A vectored Riemannian manifold is a pair (M, v), where $v \in TM$ is a unit vector, and a framed Riemannian manifold is a pair (M, f), where $f \in \mathcal{F}M$ is some (orthonormal) frame. We define

$$T^1\mathcal{M}^d = \{\text{vectored, connected, complete Riemannian } d\text{-manifolds } (M, v)\}/\sim$$
$$\mathcal{F}\mathcal{M}^d = \{\text{framed, connected, complete Riemannian } d\text{-manifolds } (M, f)\}/\sim$$

where in both cases we consider vectored (framed) manifolds up to vectored (framed) isometry. Smooth convergence of vectored Riemannian manifolds is defined as follows: we say that $(M_i, v_i) \to (M, v)$ if for every $R > 0$ there is an open set $U \supset B(p, R)$ and, for sufficiently large i, embeddings

$$f_i : U \to M_i \tag{45}$$

with $Df_i(v) = v_i$ such that and $f_i^* g_i \to g$ on U in the C^∞-topology, where g_i, g are the Riemannian metrics on M_i, M; an analogous definition gives a notion of smooth convergence on $\mathcal{F}\mathcal{M}^d$. We then have:

Theorem A.11. *$T^1\mathcal{M}^d$ and $\mathcal{F}\mathcal{M}^d$ both admit complete, separable metrics that topologize smooth convergence, and such that the natural maps*

$$\mathcal{F}\mathcal{M}^d \to T^1\mathcal{M}^d \to \mathcal{M}^d$$

defined by taking a frame to its first element, and a vector to its basepoint, are quotient maps.

The proof is identical to the work done earlier in this section. In particular, one can still reinterpret smooth convergence through locally bilipschitz maps of iterated tangent bundles, as long as these maps respect the obvious lifts of the base vectors or frames of the original manifolds. Also, all compactness arguments still apply since we have chosen unit vectors and orthonormal frames.

A.4. The Chabauty topology. Suppose M is a proper metric space and let $\mathcal{C}(M)$ be the space of closed subsets of M. The Chabauty topology on $\mathcal{C}(M)$ is that generated by subsets of the form

$$\{C \in \mathcal{C}(M) \mid C \cap K = \emptyset\}, \quad \{C \in \mathcal{C}(M) \mid C \cap U \neq \emptyset\}, \tag{46}$$

where $K \subset M$ is compact and $U \subset M$ is open. It is also called the Fell topology by analysts. Convergence can be characterized as follows:

Proposition A.12 (Prop E.12, [14])**.** *A sequence (C_i) in $\mathcal{C}(M)$ converges to $C \in \mathcal{C}(M)$ in the Chabauty topology if and only if*

1) *if $x_{i_j} \in C_{i_j}$ and $x_{i_j} \to x \in M$, where $i_j \to \infty$, then $x \in C$.*
2) *if $x \in C$, then there exist $x_i \in C_i$ such that $x_i \to x$.*

The Chabauty topology is compact, separable and metrizable [14, Lemma E.1.1]. When M is compact, it is induced by the Hausdorff metric on $\mathcal{C}(M)$, where the distance between closed subsets $C_1, C_2 \subset M$ as defined as

$$d_{Haus}(C_1, C_2) = \inf\{\epsilon \mid C_1 \subset N_\epsilon(C_2) \text{ and } C_2 \subset N_\epsilon(C_1)\}.$$

For noncompact M, the Chabauty topology is almost, but not quite, induced by taking the Hausdor topology on all compact subsets of M. Namely, x a base point p 2 M. If A M is closed and R > 0, set

$$A_R = A \setminus \overline{B(p; R)};$$

and then dene a pseudo-metric d_R on C(M) by setting

$$d_R(A; B) = \min \; \text{n} \; 1; d_{Haus}(A_R; B_R) \; \text{o} \; ;$$

where d_{Haus} is the Hausdor metric of the compact subset $\overline{B(p; R)}$ M.

The family of pseudo-metrics fd_R j R > 0g does not determine the Chabauty topology, since if x_i ! x is a convergent sequence of points with d(x; p) = R and d(x_i; p) > R for all i, then fx_ig ! fxg in the Chabauty topology, but d_R(fx_ig; fxg) = 1 for all i.

We now describe how to taper down d_R near the boundary of B(p; R) so that even when points converge into @B(p; R) from outside, d_R does not jump in the limit. The idea is to view the Hausdor distance on closed subsets of $\overline{B(p; R)}$ as a special case of a distance d_{usc} on upper semicontinuous (u.s.c.) functions

$$; g : \overline{B(p; R)} \; ! \; [0; 1]:$$

Closed subsets A; B have u.s.c. characteristic functions $1_A; 1_B$. The advantage of functions is that $1_A(x)$ and $1_B(x)$ can be scaled to converge to zero as x ! @B(p; R), so that near @B(p; R), the contribution to distance is negligible.

To dene a metric on u.s.c. functions, we use an idea of Beer [13]. Given a compact metric space K and a function : K ! [0; 1], let

$$H() = f(x; s) \text{ j } s \;\; (x)g \;\; K \;\; [0; 1]$$

be the hypograph of . The distance between functions ; : K ! [0; 1] is

$$d_{usc}(; \;\;) := d_{Haus}(H(); H(\;\;)) ;$$

where d_{Haus} is the Hausdor metric on K [0; 1], considered with the product metric d((x; s); (y; t)) = d(x; y) + d(s; t). Note that if A; B K are closed,

$$d_{usc}(1_A; 1_B) = \min 1; d_{Haus}(A; B) \;\; :$$

Fix a function : K ! [0; 1], and dene a new metric d_{Haus} on C(K) via

$$d_{Haus}(A; B) = d_{usc}(H(\;\; 1_A) ; H(\;\; 1_B)) :$$

So in words, we are just taking the Hausdor distance between A; B, but are scaling down the importance of dierent parts of the sets as dictated by .

Lemma A.13. Suppose ; : K ! [0; 1] are u.s.c. functions, that is -lipschitz, and that (x) C (x). Then d_{Haus} maxfC; + 1gd_{Haus}:

Note that the lipschitz condition is necessary for any sort of inequality. For instance, if approximates 1_{fxg} and = 1 is constant, we can make

$$1 \;\; d_{Haus}(fxg; fyg) >> d_{Haus}(fxg; fyg) \;\; 0$$

by taking y x and the approximation 1_{fxg} suciently close.

Proof. Suppose that $H(\ 1_A)$ and $H(\ 1_B)$ are -close in the Hausdor metric. We want to show $H(1_A)$ and $H(1_B)$ are (maxfC; + 1g)-close, i.e. that the two sets are each contained in (maxfC; + 1g)-neighborhoods of each other.

Let x 2 A. It suces to show that there is some (y; t) 2 $H(\ 1_B)$ with

(47) d((x; (x)); (y; t)) maxfC; + 1g:

For the same estimate will also work with (x) replaced by any s < (x), so $H(\ 1_A)$ is contained in a C-neighborhood of $H(\ 1_B)$. The proof of the other inclusion is the same, switching the roles of A; B.

We know that there is some (y; t) 2 $H(\ 1_B)$ with d((x; (x)); (y; t)) : If t = 0, then (x) , implying (x) C, and d((x; (x)); (x; 0)) C; which proves (47). So, assume t = 0. In this case, y 2 B and d(x; y) . Since is -lipschitz, j(x) (y)j . Hence,

d((x; (x)); (y; (y))) (1 +):

We now return to the problem of constructing a metric for the Chabauty topology on C(M), for a proper metric space M. Fix a point p 2 M and for each R > 0, dene a pseudometric $\hat{d}_R$ on C(M) by

(48) $\hat{d}_R = d^{R}_{Haus}$; where $_R(x) = \begin{cases} \frac{R \quad d(p;x)}{R} & d(p; x) \quad R \\ 0 & d(p; x) \quad R: \end{cases}$

Note that $\hat{d}_R$ induces the Hausdor topology on the set of compact subsets of the open ball B(p; R), but cannot tell apart subsets of M r B(p; R). Also, earlier we only dened our modied Hausdor metrics d_{Haus} for K compact, while M is not compact. However, since $_R(x) = 0$ when d(p; x) R, one can consider the above construction as taking place within K = $\overline{B(p; R)}$.

Proposition A.14. The Chabauty topology on C(M) is induced by the family of pseudo-metrics $\hat{d}_R$, for R 2 (0; 1).

By Lemma A.13, we have $\hat{d}_R$ $2\hat{d}_{R^0}$ whenever 1 R R^0, since $_R$ is 1-lipschitz and $_{R^0}$ $_R$. This implies the Chabauty topology is induced by any family $\hat{d}_{R_i}$ with R_i ! 1, although this is also clear from the proof below.

Proof. Assume that A_i ! A in the Chabauty topology. Fixing R, we want to show that the hypographs $H(_R\ 1_{A_i})$ Hausdor converge to $H(_R\ 1_A)$.

First, suppose that (x; t) 2 $H(_R\ 1_A)$. If t = 0, we have (x; t) 2 $H(_R\ 1_{A_i})$ for all i. If t = 0, then p 2 A \ B(p; R), so by Chabauty convergence, x = $\lim_i x_i$ for some sequence x_i 2 A_i. So, (x; t) is a limit of points $(x_i; t_i)$ 2 $H(_R\ 1_{A_i})$.

Next, suppose (x; t) is the limit of some sequence $(x_{i_j}; t_{i_j})$ 2 $H(_R 1_{A_{i_j}})$. Again, if t = 0 then (x; t) 2 $H(_R\ 1_A)$ automatically, so we are done. Otherwise, we can assume after passing to a further subsequence that t_{i_j} = 0 for all i_j. In this case, each x_{i_j} 2 A_{i_j}, so x = lim x_{i_j} 2 A. Hence (x; t) 2 $H(_R\ 1_A)$.

Finally, we must show that if (A_i) does not converge to A in the Chabauty topology, then there is some R with $\hat{d}_R(A_i; A)$! 0. There are two cases. Assume rst that x 2 A is not the limit of any sequence x_i 2 A_i. Taking R > d(p; x), we see that $(x;\ _R\ 1_A(x))$ 2 $H(_R\ 1_A)$ is not the limit of any sequence of points in the hypographs $H(_R\ 1_{A_i})$, so we have $\hat{d}_R(A_i; A)$! 0. Similarly, if there is some sequence x_{i_j} 2 A_{i_j} that converges to a point outside of A, the points $(x_{i_j}; {}_R\ 1_{A_{i_j}}(x_{i_j}))$ will converge to a point outside of $H({}_R\ 1_A)$.

Finally, for use in the next section, we prove:

Corollary A.15. Suppose that f : $B_{M_1}(p_1; R_1)$! M_2 is a locally -bilipschitz embedding with R_1 1, and $f(p_1) = p_2$. Then for any R_2 R_1; we have

$$\hat{d}_{R_1}(f\ ^1(C); f\ ^1(D))\ d\ \hat{}_{R_2}(C; D);\quad 8C; D\ 2\ C(M_2):$$

Note that R_2 R_1 implies f $B_{M_1}(p_1; R_1)$ $B_{M_2}(p_1; R_2)$.

Proof. The two sides of the inequality are $d^{i}_{Haus}(C; D)$, i = 1; 2, where

$$_1 : M_2\ !\ [0; 1];\ _1(x_2) = \begin{cases} \frac{R_1\ d(p_1;x_1)}{R_1} & x_2 = f(x_1);\ x_1\ 2\ B_M(p_1; R_1); \\ 0 & x_2\ 2\ f\ B_{M_1}(p_1; R_1); \end{cases}$$

$$_2 : M_2\ !\ [0; 1];\ _2(x_2) = \begin{cases} \frac{R_2\ d(p_2;x_2)}{R_2} & x_2\ 2\ B_{M_2}(p_1; R_2) \\ 0 & x_2\ 2\ B_{M_2}(p_1; R_2): \end{cases}$$

Since f is -lipschitz, we have $d(p_2; x_2)$ $d(p_1; x_1)$ if $f(x_1) = x_2$. Conversely, suppose is a path in M_2 joining x_2 to p_2. The preimage $f\ ^1()$ is either a path from x_1 to p_1, or is a union of paths, the last of which is a path from $@B_M(p_1; R_1)$ to p_1. In either case, the length of $f\ ^1()$ is at least $d(p; x_1)$, so the length of is at least $\frac{1}{}d(p; x_1)$, as f is locally $\frac{1}{}$-bilipschitz. This shows

(49) $$\frac{1}{}d(p_1; x_1)\ d(p_2; x_2)\ d(p_1; x_1):$$

Note that it may not be true that f is globally -bilipschitz, e.g. if f is the inclusion of an interval of length :999 into a circle of length 1, but (49) holds in this case because one of the two points is the center of the interval.

It follows from (49) that $_1$ is -lipschitz. Moreover, since R_2 R_1,

$$_1(x_2)\ \frac{R_1\ d(p_1; x_1)\ R_1}{R}\ \frac{d(p_2; x_2)}{R}\ _2(x_2)\ 1\ 1$$

for all x_2 2 $B_{M_2}(p_1; R_2)$. So $_1 = _2$ 1. Therefore, the hypotheses of Lemma A.13 are satised with the constant , which proves the corollary.

A.5. The smooth-Chabauty topology. In this section we combine the smooth topology of §A.2 with the Chabauty topology of §A.4. Consider the set

$$CM^d = f(M; p; C)\ j\ \begin{matrix} M \text{ a complete, connected Riemannian} \\ d\text{-manifold, } p2M; \text{ and } C\,M \text{ a closed subset} \end{matrix}\ g = ;$$

where $(M_1; p_1; C_1)$ $(M_2; p_2; C_2)$ if there is an isometry M_1 ! M_2 with p_1 ! p_2 and C_1 ! C_2. We say that $(M_i; p_i; C_i)$! $(M; p; C)$ in the smooth-Chabauty topology if for large i there are embeddings

(50) $$f_i : B_M(p; R_i)\ !\ M_i$$

with $f_i(p) = p_i$ such that R_i ! 1 and the following two conditions hold:

1) $f_i^{\ast} g_i$! g in the C^1-topology, where $g_i; g$ are the Riemannian metrics on $M_i; M$, and
2) $f_i^{\ 1}(\overline{C_i})$! C in the Chabauty topology on closed subsets of M.

Note that the metrics $f_i^* g_i$ are only partially defined, but they are defined on larger and larger subsets of M as $i \to \infty$. So as C^∞-convergence of metrics is checked on compact sets, the convergence in 1) still makes sense. As in §A.1, we call (f_i) a sequence of almost isometric maps witnessing the convergence $(M_i, p_i, C_i) \to (M, p, C)$. Also, when the R_i do not matter, we will again write

$$f_i : M \dashrightarrow M_i$$

to indicate that the maps f_i are partially defined, but that their domains of definition exhaust M. (This notation will be mostly used in the body of the paper, not in this appendix.)

We now show how to construct a quasi-metric that induces the smooth–Chabauty topology. Here, a quasi-metric on a set X is a nonnegative, symmetric function $d : X \times X \to \mathbb{R}$ that vanishes exactly on the diagonal and for some $K \geq 1$ satisfies the quasi-triangle inequality

$$d(x_1, x_3) \leq K(d(x_1, x_2) + d(x_2, x_3)), \quad \forall x_1, x_2, x_3 \in X.$$

Examples of quasi-metrics include powers $d = \delta^\alpha$ of metrics δ, and a theorem of Frink, c.f. [8], implies that for every quasi-metric d, there is an honest metric δ on X such that $\frac{1}{K}\delta^\alpha \leq d \leq K\delta^\alpha$ for some $\alpha \geq 1$ and $K > 0$. The added flexibility in the quasi-triangle inequality makes it much easier to construct quasi-metrics than metrics, and yet Frink's theorem shows that essentially, one can do as much with the former as with the latter.

Given points $X_i = (M_i, p_i, C_i) \in \mathcal{CM}^d$, $i = 1, 2$, define

$$d_{R,k}(X_1, X_2) = \min\{1, \inf\{\log\lambda + \epsilon\}\},$$

where the infimum is taken over all λ, ϵ such that there is a smooth embedding $f : B_{M_1}(p_1, R/\lambda) \to M_2$ with $f(p_1) = p_2$ such that

1) $D^k f : T^k B_{M_1}(p_1, R/\lambda) \to T^k M_2$ is locally λ-bilipschitz on the subset $Z^k_{1/\lambda}(B_{M_1}(p_1, R/\lambda)) \subset T^k M_1$ with respect to the iterated Sasaki metric,[14]
2) $\hat{d}_{R/\lambda}(C_1, f^{-1}(C_2)) \leq \epsilon$, where $\hat{d}_{R/\lambda}$ is as in Proposition A.14.

Note that if λ and λ' are at least $e = 2.718\ldots$, the 1 realizes the minimum defining $d_{R,k}$. So, everywhere below, we will always assume $\lambda, \lambda' < e$.

The functions $d_{R,k}$ are not symmetric, so later on we will symmetrize them. However, they are already 'quasi-symmetric':

Lemma A.16. If $R > 0$, $k \in \mathbb{N}$ and $X_i = (M_i, p_i, C_i) \in \mathcal{CM}^d$ for $i = 1, 2$,

$$d_{R,k}(X_2, X_1) \leq e\, d_{R,k}(X_1, X_2).$$

Proof. Suppose $d_{R,k}(X_1, X_2) < \log\lambda + \epsilon$, where the sum is that in the definition of $d_{R,k}$, and choose $f : B_{M_1}(p_1, R/\lambda) \to M_2$ as above realizing this inequality. Applying Lemma A.2, the inverse map is defined on the domain

$$f^{-1} : B_{M_2}(p_2, R/\lambda^2) \to M_1,$$

and of course is locally λ-bilipschitz. Note that the iterated derivative of f^{-1} is $(D^k f)^{-1}$, which is locally λ-bilipschitz on $D^k f(Z)$, where

$$Z = Z^k_{1/\lambda}(B_{M_1}(p_1, R/\lambda)) \subset T^k M_1.$$

[14] Here, recall that $Z^k_r(M)$ is the r-neighborhood of the zero section in $T^k(M)$.

And since $D^k f$ is locally -bilipschitz, we have

$$D^k f(Z) \quad Z_{1=^k_2}(B_{M_2}(p_2; R=^2));$$

so $D^k f \quad ^1$ is locally -bilipschitz We have

$$d_{R=^2}(C_2; \ f \quad ^1 \ ^1(C_1)) = d_{R=^2}(C_2; f(C_1))$$
$$d_{R=}(f \quad ^1(C_2); C_1)$$

where the rst inequality uses Corollary A.15. The lemma follows as $< e$:

We now prove a quasi-triangle inequality:

Lemma A.17. If $R \quad e^2; k \ 2 \ N$ and $X_i = (M_i; p_i; C_i) \ 2 \ CM^d$ for $i = 1; 2; 3$,

$$d_{R;k}(X_1; X_3) \quad e(d_{R;k}(X_1; X_2) + d_{R;k}(X_2; X_3)):$$

Proof. Suppose $d_{R;k}(X_1; X_2) < \log +$ and $d_{R;k}(X_2; X_3) < \log +$, where these sums are those in the denition of $d_{R;k}$, and choose

$$f : B_{M_1}(p_1; R=) \quad ! \quad M_2; \quad g : B_{M_2}(p_2; R=) \quad ! \quad M_3 \text{ as}$$

above realizing these inequalities. The composition

$$g \ f : B_{M_1}(p_1; R=()) \quad ! \quad M_3$$

is dened, since f is itself -lipschitz, by Corollary A.7. The iterated total derivative $D^k(g \ f) : T^k_1 B_M (p_1; R=()) \quad ! \quad T^k M_3$ is also locally -bilipschitz on the appropriate domain, so we just need to deal with condition 2). But

$$d_{R=()} \ C_1; (g \ f) \quad ^1(C_3)$$
$$d_{R=()} \ C_1; f \quad ^1(C_2) \ + d_{R=()} \ f \quad ^1(C_2); (g \ f) \quad ^1(C_3) \quad (51)$$
$$2 \, d_{R=} \ C_1; f \quad ^1(C_2) \ + d_{R=} \ C_2; g \quad ^1(C_3))$$
$$2 + \quad e(+$$
$$):$$

Here, the rst term of (51) comes from the comment after the statement of Proposition A.14, and the second term comes from Corollary A.15. This nishes the proof, since then $d_{R;k}(X_1; X_3) \quad \log() + e(+) \quad e(\log + \ + \log +) \quad e(d_{R;k}(X_1; X_2) + d_{R;k}(X_2; X_3))$.

One now proceeds as in the proof of Theoreom A.10 to construct a quasi-metric D on CM^d that induces the smooth-Chabauty topology:

$$D(X; Y) = \sum_{k=9}^{1} \frac{1}{2^k} (d_{k;k}(X; Y) + d_{k;k}(Y; X));$$

where $9 > e^2$ is chosen because of Lemma A.17. Note that the quasi-symmetry lemma (Lemma A.16) implies that symmetrizing does not change the topology induced by $d_{R;k}$. Now as mentioned above, Frink's theorem, c.f. [8], implies that there is a metric on CM^d with

$$\frac{1}{K} \ {}_D = K; \quad \text{for some} \quad 1; K > 0:$$

This allows us to prove:

Theorem A.18. $\mathcal{CM}^d$ is a Polish space.

Proof. The definition of a Cauchy sequence extends verbatim to quasi-metrics, and we claim that the quasi-metric D is complete. If $X_i = (M_i, p_i, C_i)$ is a D-Cauchy sequence, the sequence of pointed manifolds (M_i, p_i) is Cauchy for the complete metric, also called D, introduced in the proof of Theorem A.10. Hence, we can assume $(M_i, p_i) \to (M, p)$ in the smooth topology.

Fix sequences $R_n \to \infty$, $\lambda_n \to 1$ and $k_n \to \infty$. Then for each n, we have that for all $i \geq I_n$ there are maps

$$f_{n,i} : B_M(p, R_n) \longrightarrow M_i$$

with $f_{n,i}(p_1) = p_2$ such that $D^{k_n} f_{n,i}$ is locally λ_n-bilipschitz on an appropriate neighborhood of the zero section in $T^k M$. Because the Chabauty topology is compact, we can pass to a subsequence in n such that

$$\overline{f_{n,I_n}^{-1} C_{I_n}} \to C \in \mathcal{C}(M), \tag{52}$$

and passing to a further subsequence[15] we may assume that

$$\hat{d}_{R_n}(f_{n,I_n}^{-1} C_{I_n}, C) \leq \frac{1}{n} \text{ and } d_{k_n,R_n}(X_i, X_j) \leq \frac{1}{n} \text{ when } i, j \geq I_n, \tag{53}$$

where $d_{R,k}$ is as in Proposition A.17 and the second part of (53) uses that (X_i) is D-Cauchy. Now set $X = (M, p, C)$. Then for each n and $i \geq I_n$, we have

$$\begin{aligned} d_{k_n,R_n}(X, X_i) &\leq d_{k_n,R_n}(X, X_{I_n}) + d_{k_n,R_n}(X_{I_n}, X_i) \\ &\leq (\log \lambda_n + 1/n) + \frac{1}{n}. \end{aligned}$$

This converges to zero with n, so $X_i \to X$ in $\mathcal{CM}^d$. In other words, D is a complete quasi-metric. But Cauchy sequences for D are the same as Cauchy sequences for ρ, so this means that ρ is a complete metric on $\mathcal{CM}^d$.

Separability of $\mathcal{CM}^d$ follows from separability of $\mathcal{M}^d$: we can choose an element from a countable dense subset of $\mathcal{CM}^d$ by first choosing a pointed manifold (M, p) from a countable dense subset of $\mathcal{M}^d$ and then choosing a finite subset $C \subset M$ that lies within a fixed countable dense subset of M. ∎

There are a number of variants of $\mathcal{CM}^d$: one could substitute pointed manifolds with vectored or framed manifolds, or require the distinguished closed subset to lie in either the unit tangent bundle or the frame bundle. (See the space $\mathcal{P}^d_{all}$ introduced in §4.2, for instance, which is the space of pointed manifolds with distinguished closed subsets of the frame bundle.) The techniques above apply just as easily to all these situations, so we will feel free to use their metrizability without comment in the text.

[15] We are performing a bit of sleight-of-hand here in order to tame the proliferation of indices. Namely, we are passing to a subsequence in n, but then also replacing the R_n and k_n with sequences that goes to infinity more slowly. The convergence in (52) alone would not be enough to conclude the first part of (53), otherwise, for instance.

A.6. A stability result for Riemannian foliated spaces. This section is devoted to the following stability result for leaves of Riemannian foliated spaces. It is a slight strengthening of Theorem 4.3 (see also Theorem 4.1) from Lessa [68], which requires that X is compact, and implies a number of classical results on foliations, such as the Reeb local stability theorem, see [68].

Theorem A.19. *Suppose X is a d-dimensional Riemannian foliated space in which $x_i \to x$ is a convergent sequence of points. Then (L_{x_i}, x_i) is pre-compact in $\mathcal{M}^d$, and every accumulation point is a pointed Riemannian cover of (L_x, x).*

Moreover, if (M, p) is such an accumulation point, it is covered by the holonomy cover $(\tilde{L}_x^{hol}, \tilde{x})$ of (L_x, x): that is, there is are covering maps

$$(\tilde{L}_x^{hol}, \tilde{x}) \to (M, p) \to (L_x, x),$$

where the composition of the two is the holonomy covering.

Recall that the holonomy cover of a pointed leaf (L, x) is the pointed cover corresponding to the subgroup of $\pi_1(L, x)$ consisting of all loops on L with trivial holonomy. The last part of Theorem A.19 will never be used in this paper, but we think it is worth including in the theorem statement for future reference.

Before giving the proof, we note the following:

Lemma A.20. *Let X be a Riemannian foliated space and let $R > 0$. Suppose that $x_i \to x \in X$ and for each i, we have a point $y_i \in L_{x_i}$ with*

$$\limsup_{i \to \infty} d_{L_{x_i}}(x_i, y_i) \leq R.$$

Then there is a subsequence of (y_i) that converges to a point $y \in L_x$ with

$$d_{L_x}(x, y) \leq R.$$

Proof. There is some $\epsilon > 0$ with the following property: for every z in the ball $B(x, R) \subset L_x$, there is a foliated chart $\phi : \mathbb{R}^d \times T \to X$ with $\phi(0, s) = z$ and

$$d_t \left([-1, 1]^d, \mathbb{R}^d \setminus [-2, 2]^d \right) \geq \epsilon \quad \forall t \in T,$$

where d_t is the Riemannian metric on $\mathbb{R}^d$ induced from that on $\phi(\mathbb{R}^d \times \{t\})$. This is a simple consequence of the relative compactness of $B(x, R)$.

If $R < \epsilon$, then pick such a foliated chart with $\phi(0, s) = x$. For large i, we have $x_i \in \phi([-1, 1]^d \times T)$ and $d_{L_{x_i}}(x_i, y_i) < \epsilon$, so it follows that

$$y_i = \phi(a_i, t_i), \text{ where } a_i \in [-2, 2]^d, \ t_i \to s,$$

and extracting a convergent subsequence of (a_i) proves the claim.

Most likely $R \geq \epsilon$, though. In this case, construct a new sequence by choosing $z_i \in L_{x_i}$ along a geodesic from x_i to y_i so that

$$d_{L_{x_i}}(x_i, z_i) = \epsilon/2, \quad \limsup_{i \to \infty} d_{L_{x_i}}(z_i, y_i) \leq R - \epsilon/2.$$

Passing to a subsequence, (z_i) converges, so we may replace (x_i) with (z_i), thus substituting R with $R - \epsilon/2$. The proof concludes by induction.

We now prove the theorem.

Proof of Theorem A.19. Given $R > 0$ and $k \in \mathbb{N}$, the R-balls around the base points x_i within the leaves L_{x_i} have uniformly C^k-bounded geometry, in the sense of Definition A.3 in the appendix. For by compactness, this is true of the R-ball around $x \in L_x$, and for large i the coordinate charts of Denition A.3 can be transferred to L_{x_i} with arbitrarily small distortion via the vertical projections in local ow boxes, see Lemma A.20. Compare also with Lemma 4.34 of Lessa [68], in which similar arguments are used. It follows from Theorem A.5 that $(L_{x_i}; x_i)$ is precompact in M^d.

Suppose now that $x_i \to x$ in X and $(L_{x_i}; x_i) \to (M; p)$ smoothly. From the smooth convergence, it follows that for every $R > 0$ we have maps

$$f_i : B(p; R) \to L_{x_i} \subset X; \quad f(p) = x_i;$$

that are locally bilipschitz with distortion constants converging to 1. We claim:

Claim A.21. *After passing to a subsequence, the maps f_i converge to a local isometry $f : B(p; R) \to L_x$ with $f(p) = x$.*

Proof. The rst step is to construct a metric on X with respect to which the maps f_i are uniformly lipschitz, and the second is to show that the images of the f_i are contained in some compact subset of X, so that Arzela-Ascoli applies.

In [68, Lemma 4.33], Lessa shows that any compact Riemannian foliated space X admits a metric d that is adapted to the leafwise Riemannian structure: i.e. when $x; y$ lie on the same leaf, their leafwise distance is at least $d(x; y)$. The idea is to rst construct pseudo-metrics on X that vanish outside of a given foliated chart $\mathbb{R}^d \times T \to X$, by combining the leafwise metrics d_t with the distance d_T in T. Then, one covers X with a nite number of such charts, and sums the resulting pseudo-metrics to give an adapted metric on X.

In our situation, X may not be (even locally) compact, so this method fails to produce an adapted metric. However, as $B(x; R) \subset L_x$ is relatively compact, the same argument does give a pseudo metric d on X such that

1) if $x; y$ lie on the same leaf, their leafwise distance is at least $d(x; y)$,
2) there is a neighborhood $U \subset X$ of the ball $B(x; R) \subset L_x$ such that d restricts to a metric on the closure of U.

As the maps $f_i : B(p; R) \to L_{x_i} \subset X$ are locally bilipschitz with distortion constants converging to 1, they are uniformly lipschitz with respect to the adapted pseudo-metric d. Lemma A.20 implies that $f_i(B(p; R)) \subset U$ for large i and that

$$K = \overline{B(x; R)} \cup \bigcup_i \overline{f_i(B(p; R))}$$

is compact, so the intersection of $K \cap \overline{U}$ is a compact metric space that contains $f_i(B(p; R))$ for large i. By Arzela-Ascoli's theorem, f_i converges after passing to a subsequence. The limit is a local isometry $f : B(p; R) \to L_x$ with $f(p) = x$, proving Claim A.21, and therefore Theorem A.19.

Using Claim A.21, a diagonal argument now gives a local isometry

$$f : M \to L_x; \quad f(p) = x$$

dened on all of M. As M and L_x are both complete, connected Riemannian manifolds, this f is a Riemannian covering map. The fact that $\tilde{L}_x^{hol} \to L_x$ factors through M is exactly the same as in [68], see the last 3 paragraphs of the proof of Theorem 4.3.

References

1. Miklos Abert, Nicolas Bergeron, Ian Biringer, and Tsachik Gelander, Convergence of normalized betti numbers in nonpositive curvature, arXiv preprint arXiv:1811.02520 (2018).
2. Miklos Abert, Nicolas Bergeron, Ian Biringer, Tsachik Gelander, Nikolay Nikolov, Jean Raimbault, and Iddo Samet, On the growth of Betti numbers of locally symmetric spaces, C. R. Math. Acad. Sci. Paris 349 (2011), no. 15-16, 831{835. MR 2835886 (2012j:58033)
3. , On the growth of L^2-invariants for sequences of lattices in Lie groups, arXiv:1210.2961 (2012).
4. Miklos Abert, Nicolas Bergeron, and Etienne Le Masson, Eigenfunctions and random waves in the benjamini-schramm limit, arXiv preprint arXiv:1810.05601 (2018).
5. Miklos Abert, Yair Glasner, and Balint Virag, Kesten's theorem for invariant random subgroups, arXiv:1201.3399 (2012).
6. Miklos Abert and Tamas Hubai, Benjamini-Schramm convergence and the distribution of chromatic roots for sparse graphs, Combinatorica 35 (2015), no. 2, 127{151.
7. Ian Agol, Tameness of hyperbolic 3-manifolds, arXiv:math.GT/0405568 (2006).
8. Hugo Aimar, Bibiana Iaei, and Liliana Nitti, On the Macas-Segovia metrization of quasi-metric spaces, Rev. Un. Mat. Argentina 41 (1998), no. 2, 67{75. MR 1700292
9. David Aldous and Russell Lyons, Processes on unimodular random networks, Electron. J. Probab. 12 (2007), no. 54, 1454{1508. MR 2354165 (2008m:60012)
10. Jesus A. Alvarez Lopez, Ramon Barral Lijo, and Alberto Candel, A universal Riemannian foliated space, Topology Appl. 198 (2016), 47{85. MR 3433188
11. Werner Ballmann, Lectures on Kahler manifolds, ESI Lectures in Mathematics and Physics, European Mathematical Society (EMS), Zurich, 2006. MR 2243012
12. Werner Ballmann, Mikhael Gromov, and Viktor Schroeder, Manifolds of nonpositive curvature, Progress in Mathematics, vol. 61, Birkhauser Boston, Inc., Boston, MA, 1985. MR 823981 (87h:53050)
13. Gerald Beer, Upper semicontinuous functions and the Stone approximation theorem, J. Approx. Theory 34 (1982), no. 1, 1{11. MR 647707
14. Riccardo Benedetti and Carlo Petronio, Lectures on hyperbolic geometry, Universitext, Springer-Verlag, Berlin, 1992. MR MR1219310 (94e:57015)
15. Itai Benjamini and Nicolas Curien, Ergodic theory on stationary random graphs, Electron. J. Probab. 17 (2012), no. 93, 20. MR 2994841
16. Itai Benjamini, Russ Lyons, Yuval Peres, and Oded Schramm, Group-invariant percolation on graphs, Geom. Funct. Anal. 9 (1999), no. 1, 29{66. MR 1675890 (99m:60149)
17. Itai Benjamini, Russell Lyons, and Oded Schramm, Unimodular random trees, Ergodic Theory and Dynamical Systems 35 (2015), no. 2, 359{373.
18. Itai Benjamini and Oded Schramm, Recurrence of distributional limits of nite planar graphs, Electron. J. Probab. 6 (2001), no. 23, 13 pp. (electronic). MR 1873300 (2002m:82025)
19. Ian Biringer and Jean Raimbault, Ends of unimodular random manifolds, Proceedings of the American Mathematical Society 145 (2017), no. 9, 4021{4029.
20. Ian Biringer and Juan Souto, Ranks of mapping tori via the curve complex, preprint at arxiv.org/abs/1409.1878.
21. Ian Biringer and Omer Tamuz, Unimodularity of invariant random subgroups, arxiv.org/abs/1402.1042.
22. Francis Bonahon, Bouts des varietes hyperboliques de dimension 3, Ann. of Math. (2) 124 (1986), no. 1, 71{158. MR MR847953 (88c:57013)
23. Nicolas Bourbaki, Integration. II. Chapters 7{9, Elements of Mathematics (Berlin), Springer-Verlag, Berlin, 2004, Translated from the 1963 and 1969 French originals by Sterling K. Berberian. MR 2098271 (2005f:28001)
24. Lewis Bowen, Invariant random subgroups of the free group, arXiv:1204.5939 (2012).
25. , Cheeger constants and l^2-betti numbers, Duke Mathematical Journal 164 (2015), no. 3, 569{615.
26. Lewis Bowen, Rostislav Grigorchuk, and Rostyslav Kravchenko, Invariant random subgroups of lamplighter groups, Israel Journal of Mathematics 207 (2015), no. 2, 763{782.
27. Martin R. Bridson and Andre Haeiger, Metric spaces of non-positive curvature, 319 (1999), xxii+643. MR MR1744486 (2000k:53038)

28. Sara Broerio, Maura Salvatori, and Wolfgang Woess, Brownian motion and harmonic functions on Sol(p; q), Int. Math. Res. Not. IMRN (2012), no. 22, 5182{5218. MR 2997053
29. Danny Calegari and David Gabai, Shrinkwrapping and the taming of hyperbolic 3-manifolds, J. Amer. Math. Soc. 19 (2006), no. 2, 385{446 (electronic). MR MR2188131 (2006g:57030)
30. R. D. Canary, D. B. A. Epstein, and P. Green, Notes on notes of Thurston, Analytical and geometric aspects of hyperbolic space (Coventry/Durham, 1984), London Math. Soc. Lecture Note Ser., vol. 111, Cambridge Univ. Press, Cambridge, 1987, pp. 3{92. MR 903850 (89e:57008)
31. Richard D. Canary, Ends of hyperbolic 3-manifolds, J. Amer. Math. Soc. 6 (1993), no. 1, 1{35. MR MR1166330 (93e:57019)
32. , A covering theorem for hyperbolic 3-manifolds and its applications, Topology 35 (1996), no. 3, 751{778. MR MR1396777 (97e:57012)
33. Alberto Candel, The harmonic measures of Lucy Garnett, Adv. Math. 176 (2003), no. 2, 187{247. MR 1982882 (2004m:58057)
34. Alberto Candel and Lawrence Conlon, Foliations. I, Graduate Studies in Mathematics, vol. 23, American Mathematical Society, Providence, RI, 2000. MR 1732868 (2002f:57058)
35. , Foliations. II, Graduate Studies in Mathematics, vol. 60, American Mathematical Society, Providence, RI, 2003. MR 1994394 (2004e:57034)
36. Pedro J. Catuogno, Diego S. Ledesma, and Paulo R. Runo, Foliated stochastic calculus: harmonic measures, Trans. Amer. Math. Soc. 368 (2016), no. 1, 563{579. MR 3413874
37. Je Cheeger, Finiteness theorems for Riemannian manifolds, Amer. J. Math. 92 (1970), 61{74. MR 0263092 (41 #7697)
38. , Critical points of distance functions and applications to geometry, Geometric topol-ogy: recent developments (Montecatini Terme, 1990), Lecture Notes in Math., vol. 1504, Springer, Berlin, 1991, pp. 1{38. MR 1168042 (94a:53075)
39. Je Cheeger and David G. Ebin, Comparison theorems in Riemannian geometry, AMS Chelsea Publishing, Providence, RI, 2008, Revised reprint of the 1975 original. MR 2394158
40. A. Connes, A survey of foliations and operator algebras, Operator algebras and applications, Part I (Kingston, Ont., 1980), Proc. Sympos. Pure Math., vol. 38, Amer. Math. Soc., Providence, R.I., 1982, pp. 521{628. MR 679730 (84m:58140)
41. D. J. Daley and D. Vere-Jones, An introduction to the theory of point processes. Vol. I, second ed., Probability and its Applications (New York), Springer-Verlag, New York, 2003, Elementary theory and methods. MR 1950431 (2004c:60001)
42. Patrick Eberlein, Geodesic ows on negatively curved manifolds. II, Trans. Amer. Math. Soc. 178 (1973), 57{82. MR 0314084
43. G. A. Edgar, A noncompact Choquet theorem, Proc. Amer. Math. Soc. 49 (1975), 354{358. MR 0372586
44. Amichai Eisenmann and Yair Glasner, Generic IRS in free groups, after Bowen, Proceedings of the American Mathematical Society 144 (2016), no. 10, 4231{4246.
45. Jacob Feldman and Calvin C. Moore, Ergodic equivalence relations, cohomology, and von Neumann algebras. I, Trans. Amer. Math. Soc. 234 (1977), no. 2, 289{324. MR 0578656
46. Renato Feres and Abdelghani Zeghib, Leafwise holomorphic functions, Proc. Amer. Math. Soc. 131 (2003), no. 6, 1717{1725 (electronic). MR 1955258 (2003j:32040)
47. Michael Freedman, Joel Hass, and Peter Scott, Least area incompressible surfaces in 3-manifolds, Invent. Math. 71 (1983), no. 3, 609{642. MR MR695910 (85e:57012)
48. D. H. Fremlin, Measure theory. Vol. 4, Torres Fremlin, Colchester, 2006, Topological measure spaces. Part I, II, Corrected second printing of the 2003 original. MR 2462372
49. Harry Furstenberg, A note on Borel's density theorem, Proc. Amer. Math. Soc. 55 (1976), no. 1, 209{212. MR 0422497
50. Damien Gaboriau and Russell Lyons, A measurable-group-theoretic solution to von Neumann's problem, Invent. Math. 177 (2009), no. 3, 533{540. MR 2534099 (2011a:37010)
51. Lucy Garnett, Foliations, the ergodic theorem and Brownian motion, J. Funct. Anal. 51 (1983), no. 3, 285{311. MR 703080 (84j:58099)
52. , Statistical properties of foliations, Geometric dynamics (Rio de Janeiro, 1981), Lec-ture Notes in Math., vol. 1007, Springer, Berlin, 1983, pp. 294{299. MR 1691595
53. Etienne Ghys, Topologie des feuilles generiques, Ann. of Math. (2) 141 (1995), no. 2, 387{422. MR 1324140 (96b:57032)

54. Yair Glasner, Invariant random subgroups of linear groups, Israel Journal of Mathematics 219 (2017), no. 1, 215{270.
55. Gernot Greschonig and Klaus Schmidt, Ergodic decomposition of quasi-invariant probability measures, Colloq. Math. 84/85 (2000), no. part 2, 495{514, Dedicated to the memory of Anzelm Iwanik. MR 1784210
56. Misha Gromov, Metric structures for Riemannian and non-Riemannian spaces, english ed., Modern Birkhauser Classics, Birkhauser Boston, Inc., Boston, MA, 2007, Based on the 1981 French original, With appendices by M. Katz, P. Pansu and S. Semmes, Translated from the French by Sean Michael Bates. MR 2307192 (2007k:53049)
57. Olle Haggstrom, Innite clusters in dependent automorphism invariant percolation on trees, Ann. Probab. 25 (1997), no. 3, 1423{1436. MR 1457624 (98f:60207)
58. Yair Hartman and Omer Tamuz, Stabilizer rigidity in irreducible group actions, arXiv:1307.7539 (2013).
59. , Furstenberg entropy realizations for virtually free groups and lamplighter groups, Journal d'Analyse Mathematique 126 (2015), no. 1, 227{257.
60. Hamed Hatami, Lazlo Lovasz, and Balas Szegedy, Limits of locally-globally convergent graph sequences, Geom. Funct. Anal. 24 (2014), no. 1, 269{296. MR 3177383
61. Ernst Heintze and Hans-Christoph Im Hof, Geometry of horospheres, J. Dierential Geom. 12 (1977), no. 4, 481{491 (1978). MR 512919
62. Sigurdur Helgason, Dierential geometry, Lie groups, and symmetric spaces, Pure and Ap-plied Mathematics, vol. 80, Academic Press, Inc. [Harcourt Brace Jovanovich, Publishers], New York-London, 1978. MR 514561
63. V. A. Kamanovich, Brownian motion on foliations: entropy, invariant measures, mixing, Funktsional. Anal. i Prilozhen. 22 (1988), no. 4, 82{83. MR 977003
64. VA Kaimanovich, Invariance, quasi-invariance, and unimodularity for random graphs, Jour-nal of Mathematical Sciences 219 (2016), no. 5, 747{764.
65. Vadim A. Kaimanovich and Mikhail Lyubich, Conformal and harmonic measures on lami-nations associated with rational maps, Mem. Amer. Math. Soc. 173 (2005), no. 820, vi+119. MR 2111096
66. Michael Kapovich, Hyperbolic manifolds and discrete groups, Progress in Mathematics, vol. 183, Birkhauser Boston Inc., Boston, MA, 2001. MR MR1792613 (2002m:57018)
67. Anatole Katok and Vaughn Climenhaga, Lectures on surfaces, Student Mathematical Li-brary, vol. 46, American Mathematical Society, Providence, RI; Mathematics Advanced Study Semesters, University Park, PA, 2008, (Almost) everything you wanted to know about them. MR 2457970 (2010e:53004)
68. Pablo Lessa, Brownian motion on stationary random manifolds, Theses, Universite Pierre et Marie Curie - Paris VI, March 2014.
69. , Reeb stability and the Gromov-Hausdor limits of leaves in compact foliations, Asian J. Math. 19 (2015), no. 3, 433{463. MR 3361278
70. , Brownian motion on stationary random manifolds, Stoch. Dyn. 16 (2016), no. 2, 1660001, 66. MR 3470550
71. Jesu A Alvarez Lopez and Alberto Candel, On turbulent relations, arXiv preprint arXiv:1209.0307 (2012).
72. Jesu A Alvarez Lopez and Ramo Barral Lijo, Bounded geometry and leaves, arXiv preprint arXiv:1504.06835 (2015).
73. Laszlo Lovaz, Large networks and graph limits, American Mathematical Society Colloquium Publications, vol. 60, American Mathematical Society, Providence, RI, 2012. MR 3012035
74. W. Luck, Approximating L^2-invariants and homology growth, Geom. Funct. Anal. 23 (2013), no. 2, 622{663. MR 3053758
75. Russell Lyons and Yuval Peres, Probability on trees and networks, 2015.
76. Katsuhiko Matsuzaki and Masahiko Taniguchi, Hyperbolic manifolds and Kleinian groups, Oxford Mathematical Monographs, The Clarendon Press Oxford University Press, New York, 1998, Oxford Science Publications. MR MR1638795 (99g:30055)
77. Curtis T. McMullen, Renormalization and 3-manifolds which ber over the circle, An-nals of Mathematics Studies, vol. 142, Princeton University Press, Princeton, NJ, 1996. MR MR1401347 (97f:57022)
78. Kam Ping Mok, On the dierential geometry of frame bundles of Riemannian manifolds, J. Reine Angew. Math. 302 (1978), 16{31. MR 511689 (80c:53057)

79. Calvin C. Moore and Claude L. Schochet, Global analysis on foliated spaces, second ed., Mathematical Sciences Research Institute Publications, vol. 9, Cambridge University Press, New York, 2006. MR 2202625 (2006i:58035)
80. Hossein Namazi, Pekka Panka, and Juan Souto, Distributional limits of Riemannian manifolds and graphs with sublinear genus growth, arxiv.org/abs/1208.2130 (Preprint 2012).
81. Peter Petersen, Riemannian geometry, second ed., Graduate Texts in Mathematics, vol. 171, Springer, New York, 2006. MR 2243772 (2007a:53001)
82. Robert R. Phelps, Lectures on Choquet's theorem, second ed., Lecture Notes in Mathematics, vol. 1757, Springer-Verlag, Berlin, 2001. MR 1835574
83. John G. Ratclie, Foundations of hyperbolic manifolds, Graduate Texts in Mathematics, vol. 149, Springer-Verlag, New York, 1994. MR 1299730
84. Ludovic Riord, A Morse-Sard theorem for the distance function on Riemannian manifolds, Manuscripta Math. 113 (2004), no. 2, 251{265. MR 2128549 (2005m:53059)
85. V. A. Rohlin, On the fundamental ideas of measure theory, Amer. Math. Soc. Translation 1952 (1952), no. 71, 55. MR 0047744
86. Shigeo Sasaki, On the dierential geometry of tangent bundles of Riemannian manifolds, Tôhoku Math. J. (2) 10 (1958), 338{354. MR 0112152 (22 #3007)
87. Michael Schrödl, ‘^{2}-betti numbers of random rooted simplicial complexes, arXiv preprint arXiv:1810.09716 (2018).
88. David Simmons, Conditional measures and conditional expectation; Rohlin's disintegration theorem, Discrete Contin. Dyn. Syst. 32 (2012), no. 7, 2565{2582. MR 2900561
89. Simon Thomas and Robin Tucker-Drob, Invariant random subgroups of strictly diagonal limits of nite symmetric groups, arXiv preprint arXiv:1402.4837 (2014).
90. William Thurston, The geometry and topology of three-manifolds, Lecture notes at Princeton University (1980).
91. William P. Thurston, Hyperbolic structures on 3-manifolds. I. Deformation of acylindrical manifolds, Ann. of Math. (2) 124 (1986), no. 2, 203{246. MR MR855294 (88g:57014)
92. V.S. Varadarajan, Groups of automorphisms of Borel spaces, Transactions of the American Mathematical Society 109 (1963), no. 2, 191{220.
93. Friedhelm Waldhausen, On irreducible 3-manifolds which are suciently large, Ann. of Math. (2) 87 (1968), 56{88. MR MR0224099 (36 #7146)
94. F. W. Warner, Extensions of the Rauch comparison theorem to submanifolds, Trans. Amer. Math. Soc. 122 (1966), 341{356. MR 0200873

Miklos Abert, Renyi Institute of Mathematics, 13-15 Realtanoda utca, 1053 Budapest, Hungary,

Email address: abert.miklos@renyi.mta.hu

Ian Biringer, Boston College, Mathematics Department, 140 Commonwealth Ave. Chestnut Hill, MA 02467-3806,

Email address: ianbiringer@gmail.com