msp

Geometry & Topology 22 (2018) 2757–2790

Rigidity for convex-cocompact actions on rank-one symmetric spaces

GUY C DAVID
KYLE KINNEBERG

When $\Gamma \curvearrowright X$ is a convex-cocompact action of a discrete group on a noncompact rank-one symmetric space X, there is a natural lower bound for the Hausdorff dimension of the limit set $\Lambda(\Gamma) \subset \partial X$, given by the Ahlfors regular conformal dimension of $\partial\Gamma$. We show that equality is achieved precisely when Γ stabilizes an isometric copy of some noncompact rank-one symmetric space in X on which it acts with compact quotient. This generalizes a theorem of Bonk and Kleiner, who proved it in the case that X is real hyperbolic.

To prove our main theorem, we study tangents of Lipschitz differentiability spaces that are embedded in a Carnot group $\mathbb{G}$. We show that almost all tangents are isometric to a Carnot subgroup of $\mathbb{G}$, at least when they are rectifiably connected. This extends a theorem of Cheeger, who proved it for PI spaces that are embedded in Euclidean space.

53C24, 53C35; 53C17, 53C23

1 Introduction

A classic theorem of Bowen [12] states that the limit set of a convex-cocompact action of a Fuchsian group on $\mathbb{H}^3_{\mathbb{R}}$ has Hausdorff dimension at least 1, and equality holds precisely when the limit set is a round circle in $\mathbb{S}^2 = \partial\mathbb{H}^3_{\mathbb{R}}$. In this case, the convex-cocompact action stabilizes an isometrically embedded copy of $\mathbb{H}^2_{\mathbb{R}}$ in $\mathbb{H}^3_{\mathbb{R}}$. Expressed another way, this result means that among all quasi-Fuchsian representations of a Fuchsian group, the Hausdorff dimension of the limit set attains its minimum precisely on the Fuchsian locus.

Bowen's theorem was later generalized to convex-cocompact actions of Fuchsian groups on any $\mathrm{CAT}(-1)$ metric space by Bonk and Kleiner [7, Theorem 1.1]. This answered a question of Bourdon, who had proven that the analogous result holds for any uniform lattice in a noncompact rank-one symmetric space $S \neq \mathbb{H}^2_{\mathbb{R}}$ [10, Theorem 0.3]. To

Published: 1 June 2018 DOI: 10.2140/gt.2018.22.2757

state their results precisely, we need some notation. If X is a $\mathrm{CAT}(-1)$ space, then its visual boundary ∂X admits a natural class of Möbius equivalent metrics

$$d_p(x, y) = e^{-(x,y)_p},$$

where $p \in X$ and $(x, y)_p$ denotes the Gromov product based at p. In particular, these metrics are bi-Lipschitz equivalent. The Hausdorff dimension of ∂X, and also of any subset, is therefore well-defined. If $\Gamma \curvearrowright X$ is a convex-cocompact action by a discrete group Γ, then the limit set $\Lambda(\Gamma)$ of the action is a subset of ∂X that is invariant under the induced boundary action on ∂X.

The following statement puts together the results of Bourdon (all cases $S \neq \mathbb{H}^2_{\mathbb{R}}$) and Bonk and Kleiner (the $S = \mathbb{H}^2_{\mathbb{R}}$ case).

Theorem 1.1 [10; 7] *Let Γ be a uniform lattice in a noncompact rank-one symmetric space S, and suppose that $\Gamma \curvearrowright X$ is a convex-cocompact action on a $\mathrm{CAT}(-1)$ space X. Then*

$$\mathrm{Hdim}(\Lambda(\Gamma)) \geq \mathrm{Hdim}(\partial S),$$

and equality holds if and only if the action stabilizes an isometrically embedded copy of S in X on which Γ acts with compact quotient.

We should note that the case of equality is also characterized by Möbius equivalence of the limit set $\Lambda(\Gamma)$ and ∂S [10, Theorem 0.1].

Remark 1.2 The notation $\Lambda(\Gamma)$ is slightly misleading, as the limit set depends on the action of Γ on X, not just on the group Γ. More accurately, one should think of the action as a representation of Γ into the isometry group $\mathrm{Isom}(X)$, so the limit set is really the limit set of this representation. In many important cases, the space of convex-compact representations forms a moduli space with a rich geometric structure, and one can ask interesting questions about how the limit sets change as the representations vary (see eg Bridgeman, Canary, Labourie and Sambarino [13, Corollary 1.8]). However, as we focus only on a single action/representation at a time, we do not make this distinction explicit, and we retain the notation $\Lambda(\Gamma)$.

More recently, there has been interest in knowing whether similar types of rigidity phenomena occur when Γ is not assumed, a priori, to be a lattice. (For example, see work of Bonk and Kleiner [6] and the second author [35].) The first difficulty encountered is to determine a natural lower bound for $\mathrm{Hdim}(\Lambda(\Gamma))$. For example, free groups admit convex-cocompact actions on any noncompact rank-one symmetric

space with $\Lambda(\Gamma)$ equal to a Cantor set of arbitrarily small Hausdorff dimension. More explicitly, choosing loxodromic elements $g_1, \dots, g_k \in \mathrm{Isom}(S)$ with pairwise disjoint fixed points, for large enough n the subgroup generated by $g_1^n, \dots, g_k^n$ will be a free group of rank k, and the Hausdorff dimension of the limit set will tend to 0 as n goes to infinity.

In the case of Theorem 1.1, the inequality comes from work of Pansu [42, Theorem 5.5]. In standard formulation, Pansu's result says that the Hausdorff dimension of any metric space that is quasisymmetrically equivalent to ∂S must be at least as large as $\mathrm{Hdim}(\partial S)$. Stated another way, this means that whenever Γ is (virtually) a uniform lattice in S, the Ahlfors regular conformal dimension $\mathrm{confdim}_{\mathrm{AR}}(\partial\Gamma)$ of its Gromov boundary is equal to $\mathrm{Hdim}(\partial S)$. On the other hand, we note that $\mathrm{confdim}_{\mathrm{AR}}(\partial\Gamma) = 0$ whenever Γ is a free group.

It makes sense, then, that the appropriate generalization of Theorem 1.1 to Gromov hyperbolic groups would make use of the conformal dimension. In a different paper [8], Bonk and Kleiner established the following result, which concerns actions of hyperbolic groups on real hyperbolic space.

Theorem 1.3 [8, Theorem 1.4] *Let Γ be a nonelementary Gromov hyperbolic group, and let $Q = \mathrm{confdim}_{\mathrm{AR}}(\partial\Gamma)$. If $\Gamma \curvearrowright \mathbb{H}^{n+1}_{\mathbb{R}}$ is a convex-cocompact action, with $n \geq 1$, then*

$$\mathrm{Hdim}(\Lambda(\Gamma)) \geq Q,$$

and equality holds if and only if $Q = k \geq 1$ is an integer and Γ stabilizes an isometric copy of $\mathbb{H}^{k+1}_{\mathbb{R}}$ in $\mathbb{H}^{n+1}_{\mathbb{R}}$ on which it acts with compact quotient.

Our statement of this theorem is slightly different from that given in [8], but they are easily seen to be equivalent. We make this clear in Remark 3.4.

The main goal of the present paper is to extend Theorem 1.3 to actions on any noncompact rank-one symmetric space.

Theorem 1.4 *Let Γ be a nonelementary Gromov hyperbolic group, and let $Q = \mathrm{confdim}_{\mathrm{AR}}(\partial\Gamma)$. If $\Gamma \curvearrowright X$ is a convex-cocompact action on a noncompact rank-one symmetric space X, then*

$$\mathrm{Hdim}(\Lambda(\Gamma)) \geq Q,$$

and equality holds if and only if Γ stabilizes an isometric copy of a noncompact rank-one symmetric space S in X on which it acts with compact quotient.

In the equality case, ∂S and $\Lambda(\Gamma)$ are Möbius equivalent, $Q = \mathrm{Hdim}(\partial S)$ is necessarily an integer, and the symmetric space S is determined uniquely by Q and the topological dimension of $\partial\Gamma$. Indeed, the noncompact rank-one symmetric spaces are uniquely determined by the Hausdorff and topological dimensions of their visual boundaries (see Mackay and Tyson [40, page 34]). Moreover, Γ is a finite extension of a uniform lattice in $\mathrm{Isom}(S)$.

Our proof of Theorem 1.4 follows the basic outline that Bonk and Kleiner use to prove Theorem 1.3. First, the inequality comes straight from the general theory of boundaries of hyperbolic groups. Appealing to the prior work of Bonk and Kleiner [7] also dispenses with the equality case when $Q \leq 1$. When $Q > 1$ and equality holds, Bonk and Kleiner show that the limit set $\Lambda(\Gamma) \subset \mathbb{S}^n$ supports a Poincaré inequality (in the sense of Heinonen and Koskela [27]). As $\Lambda(\Gamma)$ is isometrically embedded in $\mathbb{R}^{n+1}$, a theorem of Cheeger [16] guarantees that $\Lambda(\Gamma)$ has a tangent that is isometric to some Euclidean space $\mathbb{R}^k$. This implies that $\Lambda(\Gamma)$ is bi-Lipschitz equivalent to $\mathbb{S}^k$, and they can use [6, Theorem 1.1] to conclude.

In the setting of Theorem 1.4, when $Q > 1$ and equality holds, the work of Bonk–Kleiner still implies that the limit set $\Lambda(\Gamma)$ supports a Poincaré inequality. However, it may not be embedded in any Euclidean space, so we cannot directly apply Cheeger's theorem. Instead, $\Lambda(\Gamma)$ is locally embedded in some Carnot group, namely, the Carnot group that models the local geometry of ∂X. The main work we do is to generalize Cheeger's result to spaces embedded in Carnot groups, showing that their tangents are isometric to Carnot subgroups. This is the content of the following theorem, which may be of independent interest. This result is stated and proven in a more general context as Theorem 4.1 below.

Theorem 1.5 *Let $\mathbb{G}$ be a Carnot group with homogeneous distance d. Let $X \subset \mathbb{G}$ be a closed subset such that there is a Radon measure μ on X for which (X, d, μ) is a doubling metric measure space supporting a Poincaré inequality. Let (U, ϕ) be a k–dimensional differentiability chart for X.*

Then, for μ–almost every $x \in U$, the set $\mathrm{Tan}(X, x)$ consists of one element, which is canonically isometric to a Carnot subgroup of $\mathbb{G}$ generated by a fixed k–dimensional subspace of the horizontal layer of $\mathbb{G}$.

All the terminology in Theorem 1.5 is introduced in Section 2.

The remainder of the proof of Theorem 1.4 follows similar ideas as in the endgame of [6, Theorem 1.2], though we will need an additional input that identifies, among all Carnot groups, those that locally model boundaries of rank-one symmetric spaces. This is provided by recent work of Cowling and Ottazzi [22].

The structure of the paper is as follows. Section 2 contains basic definitions and notation. Section 3 contains background material on boundaries of symmetric spaces and Gromov hyperbolic groups, as well as an analog of a result of Connell [20] that we prove as a short application of our methods (Proposition 3.6). Section 4 contains the proof of Theorem 4.1 about tangents of spaces embedded in Carnot groups, and describes some applications to bi-Lipschitz nonembedding. Finally, Section 5 contains the proof of Theorem 1.4.

Acknowledgments David was partially supported by the National Science Foundation under Grants No. DMS-1664369 and DMS-1758709.

2 Preliminaries

If (X, d) is a metric space, we write $B(x, r)$ for the open ball of radius r centered at $x \in X$ and $\overline{B}(x, r)$ for the closed ball. A metric space (X, d) is said to be *metrically doubling* if there is a constant C such that every ball of radius $r > 0$ in X can be covered by at most C balls of radius $\frac{r}{2}$.

A *metric measure space* (X, d, μ) is a complete, separable metric space (X, d) equipped with a Radon measure μ. We say that the measure μ is *doubling* if there is a constant C such that

$$\mu(B(x, 2r)) \leq C\mu(B(x, r))$$

for every ball $B(x, r)$ in X. This readily implies that (X, d) is metrically doubling [26].

For $Q > 0$, we will often consider the Q–dimensional Hausdorff measure $\mathcal{H}^Q$ on a metric space (X, d). We write $\mathrm{Hdim}(X)$ for the Hausdorff dimension of X. A metric space is *Ahlfors Q–regular*, for $Q > 0$, if there is a constant $C > 0$ such that

$$C^{-1} r^Q \leq \mathcal{H}^Q(\overline{B}(x, r)) \leq C r^Q$$

for all $x \in X$ and $0 < r \leq \mathrm{diam}(X)$. In this case, $(X, d, \mathcal{H}^Q)$ is easily seen to be a doubling metric measure space. It is also uniformly perfect [40, pages 13–14].

The *cross-ratio* of a quadruple of distinct points (x_1, x_2, x_3, x_4) in X is the value

$$[x_1, x_2, x_3, x_4] = \frac{d(x_1, x_3)d(x_2, x_4)}{d(x_1, x_4)d(x_2, x_3)}.$$

A homeomorphism $f\colon X \to Y$ between two metric spaces is *Möbius* if it preserves the cross-ratio, ie if

$$[f(x_1), f(x_2), f(x_3), f(x_4)] = [x_1, x_2, x_3, x_4]$$

for every quadruple (x_1, x_2, x_3, x_4) of distinct points in X.

More generally, given a homeomorphism $\eta\colon [0,\infty) \to [0,\infty)$, we say that a homeomorphism $f\colon X \to Y$ is *η–quasi-Möbius* (or simply *quasi-Möbius*) if

$$[f(x_1), f(x_2), f(x_3), f(x_4)] \le \eta([x_1, x_2, x_3, x_4])$$

for every quadruple (x_1, x_2, x_3, x_4) of distinct points in X. Similarly, f is said to be *η–quasisymmetric* (or simply *quasisymmetric*) if

$$\frac{d(f(x_1), f(x_2))}{d(f(x_1), f(x_3))} \le \eta\left(\frac{d(x_1, x_2)}{d(x_1, x_3)}\right)$$

for every triple (x_1, x_2, x_3) of distinct points in X.

These definitions also make sense if $f\colon X \to Y$ is assumed only to be injective, in which case it is necessarily a homeomorphism onto its image. Such a map f is called a Möbius embedding, quasi-Möbius embedding, etc. It is not difficult to show that every η–quasisymmetric mapping is $\widetilde{\eta}$–quasi-Möbius, where $\widetilde{\eta}$ depends only on η. If X and Y are bounded, then, conversely, every quasi-Möbius homeomorphism is quasisymmetric [48].

In the case that X and Y are bounded and η is linear, each η–quasi-Möbius map $f\colon X \to Y$ is actually bi-Lipschitz [35, Remark 3.2]. In particular, every Möbius map is bi-Lipschitz.

The *Ahlfors regular conformal dimension* of a doubling and uniformly perfect metric space X, denoted by

$$\operatorname{confdim}_{\mathrm{AR}}(X),$$

is the infimal Hausdorff dimension among all Ahlfors regular metric spaces quasisymmetrically homeomorphic to X. The doubling and uniformly perfect properties guarantee that $\operatorname{confdim}_{\mathrm{AR}}(X)$ is finite [26, Corollary 14.15], and it is obviously quasisymmetrically invariant. However, the infimum is not always achieved, and it is a

difficult problem to find general conditions under which it is. For more background on this quantity, we refer the reader to [8], where it plays a crucial role.

2.1 Tangents

We now briefly describe the notion of a "tangent" of a metric space, which relies on the framework of pointed Gromov–Hausdorff convergence of spaces and mappings. For the necessary background on this framework, we refer the reader to any of [32, Section 5; 36, Section 3.2; 23, Section 2].

A *pointed metric space* (X, d, x) is simply a metric space (X, d) together with a fixed point $x \in X$. If the metric d is clear, we sometimes suppress it and simply write (X, x).

Suppose that (X, d) is a doubling metric space. Fix $x \in X$ and let $\lambda_j \to \infty$ be a sequence of positive scaling factors. The sequence of pointed metric spaces $(X, \lambda_j d, x)$ then has a subsequence that converges in the pointed Gromov–Hausdorff sense to a complete pointed metric space $(\widehat{X}, \widehat{x})$. We say that $(\widehat{X}, \widehat{x})$ is a *tangent* of X at x. The set of pointed isometry classes of tangents at x is denoted by $\operatorname{Tan}(X, x)$.

If $f\colon X \to Y$ is a Lipschitz map taking values in a complete, doubling metric space (Y, ρ) with $y = f(x)$, then the sequence of Lipschitz maps

$$f\colon (X, \lambda_j d, x) \to (Y, \lambda_j \rho, y)$$

has a subsequence that converges to a Lipschitz map

$$\hat{f}\colon (\widehat{X}, \widehat{x}) \to (\widehat{Y}, \widehat{y}),$$

where $(\widehat{Y}, \widehat{y}) \in \operatorname{Tan}(Y, y)$. We will repeatedly use the "hat" notation for tangents and tangent maps. Note that if we first fix $(\widehat{X}, \widehat{x}) \in \operatorname{Tan}(X, x)$, then by passing to further subsequences, we may find a corresponding $(\widehat{Y}, \widehat{y}) \in \operatorname{Tan}(Y, y)$ and tangent map $\hat{f}\colon \widehat{X} \to \widehat{Y}$. The set of tangents and mappings obtained in this way, up to isometric equivalence, is denoted by $\operatorname{Tan}(X, x, f)$.

Remark 2.1 When $Y = \mathbb{R}^n$ is a finite-dimensional Euclidean space, we always identify $(\widehat{Y}, \widehat{y})$ with $(\mathbb{R}^n, 0)$ in the canonical way, so that $\hat{f}(x) = 0$. In this way, tangent maps behave like derivatives in the classical sense.

2.2 Lipschitz differentiability spaces

In the seminal paper [16], Cheeger introduced a type of differentiable structure for real-valued Lipschitz functions on metric measure spaces. He showed that a large class

of metric measure spaces, the so-called "PI spaces" (those which are doubling and have a Poincaré inequality in the sense of [27]), support such a structure.

The following definition is due to Cheeger [16].

Definition 2.2 A metric measure space (X, d, μ) is called a *Lipschitz differentiability space* if it satisfies the following condition. There are countably many Borel sets ("charts") U_i covering X, positive integers n_i (the "dimensions of the charts") and Lipschitz maps $\phi_i\colon X \to \mathbb{R}^{n_i}$ with respect to which any Lipschitz function $f\colon X \to \mathbb{R}$ is differentiable almost everywhere, in the sense that for each i and for μ–almost every $x \in U_i$, there exists a unique $\nabla f(x) \in \mathbb{R}^{n_i}$ such that

$$\lim_{y \to x} \frac{|f(y) - f(x) - \nabla f(x) \cdot (\phi_i(y) - \phi_i(x))|}{d(x, y)} = 0. \tag{2-1}$$

Here $\nabla f(x) \cdot (\phi_i(y) - \phi_i(x))$ denotes the standard scalar product in $\mathbb{R}^{n_i}$.

Note that the Borel measurability of the function $x \mapsto \nabla f(x)$ and the set of differentiability points of f are consequences of this definition; see [4, Remark 1.2]. In recent years, the study of Lipschitz differentiability spaces in their own right has become an active area of research, and we refer the reader to [1; 2; 3; 16; 17; 18; 23; 32; 36] for more background.

If (X, d) is doubling, it is natural to rephrase (2-1) in terms of tangents. Letting $df(x) \in \mathrm{Hom}(\mathbb{R}^n, \mathbb{R})$ be the dual functional to $\nabla f(x)$, the asymptotic in (2-1) is equivalent to the following condition. If $(\widehat{X}, \hat{x}) \in \mathrm{Tan}(X, x)$ and we have associated tangent mappings $\hat{f}\colon \widehat{X} \to \mathbb{R}$ and $\widehat{\phi}\colon \widehat{X} \to \mathbb{R}^n$ to f and ϕ, respectively, then $\hat{f}$ factors through $\mathbb{R}^n$ via $\hat{f} = df(x) \circ \widehat{\phi}$.

A similar factorization works for any Lipschitz map $f\colon X \to \mathbb{R}^N$ into a Euclidean space, simply by considering the coordinate functions for f. Namely, if (U, ϕ) is an n–dimensional chart in X, then for almost every $x \in U$, there is a unique linear map $Df(x) \in \mathrm{Hom}(\mathbb{R}^n, \mathbb{R}^N)$ such that for every $(\widehat{X}, \hat{x}) \in \mathrm{Tan}(X, x)$, any tangent map $\hat{f}\colon \widehat{X} \to \mathbb{R}^N$ factors through $\mathbb{R}^n$ via

$$\hat{f} = Df(x) \circ \widehat{\phi}.$$

Not surprisingly, the rows of the matrix for $Df(x)$ are simply the gradients $\nabla f_i(x)$ of the coordinate functions $f = (f_1, \ldots, f_N)$.

Recently, there has been much interest in understanding the geometric structure of tangents of Lipschitz differentiability spaces. One result in this vein is due to Cheeger,

Kleiner and Schioppa [18], and it will be important for us below. First, another piece of terminology.

A Lipschitz map $f\colon X \to Y$ between two metric spaces is called a *Lipschitz quotient* if it is Lipschitz and there exists $c > 0$ such that

$$B(f(x), cr) \subset f(B(x, r))$$

for all $x \in X$ and $r > 0$. The supremum of such constants c is called the *co-Lipschitz* constant of f. Note that a Lipschitz quotient map is clearly surjective.

We say that f is a *metric submersion* if it is both 1–Lipschitz and has co-Lipschitz constant $c = 1$.

Proposition 2.3 [18, Theorem 1.12] *Let (X, d) be a complete, doubling metric space, and let μ be a Radon measure on X such that (X, d, μ) is a Lipschitz differentiability space. Let (U, ϕ) be an n–dimensional chart in X.*

Then for μ–almost every $x \in U$, there is a norm $\|\cdot\|_x$ on $\mathbb{R}^n$ with the following property. For every $(\widehat{X}, \hat{x}) \in \mathrm{Tan}(X, x)$ with tangent map $\widehat{\phi}\colon \widehat{X} \to \mathbb{R}^n$, the mapping $\widehat{\phi}$ is a metric submersion onto $(\mathbb{R}^n, \|\cdot\|_x)$.

In particular, the mapping $\widehat{\phi}$ is a Lipschitz quotient onto $\mathbb{R}^n$ with its standard Euclidean metric.

The statement about Lipschitz quotients in Proposition 2.3 already appeared in earlier (independent) work of Schioppa [44, Theorem 5.56; 45] and the first author [23, Corollary 5.1]. This statement alone is enough for the application to our main result, Theorem 1.4. We state the stronger version about metric submersions only for the applications in Section 3.2 below.

As Cheeger observed, a consequence of the surjectivity of $\widehat{\phi}$ provided by Proposition 2.3 is the following result:

Corollary 2.4 *Let (U, ϕ) be an n–dimensional chart in a complete, metrically doubling Lipschitz differentiability space (X, d, μ), and suppose that $F\colon X \to \mathbb{R}^N$ is a bi-Lipschitz embedding. Then for μ–almost every $x \in U$, the set $\mathrm{Tan}(F(X), F(x))$ consists of one element, canonically isometric to a fixed n–dimensional linear subspace of $\mathbb{R}^N$.*

For a set $E \subset \mathbb{R}^N$ and a point $x \in E$, to say that each element of $\mathrm{Tan}(E,x)$ is *canonically isometric* to $\widehat{E} \subset \mathbb{R}^N$ means that, whenever $\lambda_j \to \infty$ and

$$(\lambda_j(E-x),0)$$

converges in the pointed Hausdorff sense in $\mathbb{R}^N$, the limit is $(\widehat{E},0)$. This is stronger than saying that $\mathrm{Tan}(E,x)$ consists only of one element, and hence Cheeger [16, Section 14] uses the phrase "unique in the strong sense" for this.

Proof of Corollary 2.4 We give a brief sketch of the proof, whose details can be found in Theorem 14.1 of [16] or Corollary 8.1 of [23].

Let $Y = F(X) \subset \mathbb{R}^N$. Consider a point $x \in X$ at which Proposition 2.3 applies, and at which F is differentiable with derivative $DF(x)\colon \mathbb{R}^n \to \mathbb{R}^N$. Let $y = F(x)$ and consider any tangent $(\widehat{Y},\widehat{y}) \in \mathrm{Tan}(Y,y)$, which we can view as a pointed Hausdorff limit of rescalings of $Y-y$ in $\mathbb{R}^N$. By passing to subsequences, we may obtain an associated tangent $(\widehat{X},\widehat{x}) \in \mathrm{Tan}(X,x)$ along with associated tangent mappings $\widehat{\phi}\colon \widehat{X} \to \mathbb{R}^n$ and $\widehat{F}\colon \widehat{X} \to \widehat{Y} \subset \mathbb{R}^N$. Note that $\widehat{F}$ is a bi-Lipschitz homeomorphism onto $\widehat{Y}$. In addition, $\widehat{\phi}$ is surjective by Proposition 2.3.

Differentiability implies that $\widehat{F}$ factors as $DF(x)\circ\widehat{\phi}$. Since $\widehat{F}$ is bi-Lipschitz, $DF(x)$ must be injective. It follows from this and the surjectivity of $\widehat{\phi}$ that

$$\widehat{Y} = F(\widehat{X}) = DF(x)\circ\widehat{\phi}(\widehat{X}) = DF(x)(\mathbb{R}^n),$$

which is a fixed n–dimensional linear subspace of $\mathbb{R}^N$. □

Corollary 2.4 is precisely the statement used in the Bonk–Kleiner proof of Theorem 1.3. As discussed above, it is this result that we must generalize to Carnot group targets.

2.3 Carnot groups

Here we give some brief background on Carnot groups. For more, we refer the reader to [41] or [15, Section 2].

A Carnot group is a simply connected nilpotent Lie group $\mathbb{G}$ whose Lie algebra $\mathfrak{g}$ admits a stratification

$$\mathfrak{g} = V_1 \oplus \cdots \oplus V_s,$$

where the first layer V_1 generates the rest via $V_{i+1} = [V_1, V_i]$ for all $1 \le i \le s$, and we set $V_{s+1} = \{0\}$. The exponential map $\exp\colon \mathfrak{g} \to \mathbb{G}$ is a diffeomorphism, so choosing a basis for $\mathfrak{g}$ gives exponential coordinates for $\mathbb{G}$. Equipped with any left-invariant

Riemannian metric, $\mathbb{G}$ is complete and the Lie exponential map coincides with the Riemannian exponential map. For each $x \in \mathbb{G}$, we will use $L_x\colon \mathbb{G} \to \mathbb{G}$ to denote the left multiplication map $y \mapsto xy$.

A natural family of automorphisms of $\mathbb{G}$ are the dilations $\delta_\lambda\colon \mathbb{G} \to \mathbb{G}$ for $\lambda > 0$. On the Lie algebra level, these are defined by

$$v \mapsto \lambda^i v \quad \text{for } v \in V_i, \tag{2-2}$$

and one can see that this gives a Lie algebra isomorphism. Conjugating this back to $\mathbb{G}$ by the exponential map gives δ_λ.

On any Carnot group $\mathbb{G}$, there are metrics that interact nicely with the translations and dilations, in the sense that each L_x is an isometry and δ_λ scales distances by the factor λ.

Definition 2.5 A metric $d\colon \mathbb{G} \times \mathbb{G} \to \mathbb{R}_{\geq 0}$ is called a *homogeneous distance* if it induces the manifold topology of $\mathbb{G}$, it is left-invariant, ie

$$d(xy, xz) = d(y, z) \quad \text{for all } x, y, z \in \mathbb{G},$$

and it is 1–homogeneous with respect to the dilations δ_λ defined above:

$$d(\delta_\lambda(x), \delta_\lambda(y)) = \lambda d(x, y) \quad \text{for all } \lambda > 0 \text{ and } x, y \in \mathbb{G}.$$

For example, given an inner product on the horizontal layer V_1, the associated (sub-Riemannian) Carnot–Carathéodory metric d_{CC} is a homogeneous distance.

Remark 2.6 It is a simple fact that, for any two homogeneous distances d_1, d_2 on $\mathbb{G}$, the identity map $(\mathbb{G}, d_1) \to (\mathbb{G}, d_2)$ is a bi-Lipschitz homeomorphism. In particular, a curve is rectifiable in $(\mathbb{G}, d_1)$ if and only if it is rectifiable in $(\mathbb{G}, d_2)$.

Let d be a homogeneous distance on $\mathbb{G}$, and let μ be the associated Hausdorff measure, so that $(\mathbb{G}, d, \mu)$ is a metric measure space. The left-invariance and 1–homogeneity of d easily imply that $(\mathbb{G}, d, \mu)$ is Ahlfors regular. Another consequence is that every tangent to any point in $(\mathbb{G}, d)$ is isometric to $(\mathbb{G}, d, 0)$ itself, much like tangents of finite-dimensional Euclidean spaces. As we did for those, we will identify any $(\widehat{\mathbb{G}}, \widehat{x}) \in \mathrm{Tan}(\mathbb{G}, x)$ with $(\mathbb{G}, 0)$. Similarly, if E is a closed subset of a Carnot group $\mathbb{G}$, then any tangent of E can be viewed as a pointed Hausdorff limit of pointed rescalings of E in $\mathbb{G}$ and identified with a pointed closed subset $(\widehat{E}, 0)$ of $\mathbb{G}$. The

terminology, used in $\mathbb{R}^n$ in Corollary 2.4, that tangents can be "canonically isometric" to subsets of $\mathbb{G}$ therefore also makes sense in Carnot groups. Namely, we say that each element of $\mathrm{Tan}(E,x)$ is canonically isometric to $\widehat{E}\subset\mathbb{G}$ if, whenever $\lambda_j\to\infty$ and

$$\big(\delta_{\lambda_j}(L_{-x}(E)),0\big)$$

converges in the pointed Hausdorff sense, the limit is $(\widehat{E},0)$.

For the remainder of this section, we equip $\mathbb{G}$ with a fixed homogeneous distance d and its corresponding Hausdorff measure. We can then talk about Lipschitz functions on $\mathbb{G}$ or Lipschitz mappings between Carnot groups. Note that the collection of such maps does not depend on the chosen homogeneous distance.

Lipschitz mappings between Carnot groups admit a form of differentiation more robust than that discussed in the previous subsection, one which takes into account the group structure.

Theorem 2.7 (Pansu [43]) *Let $f\colon\mathbb{G}_1\to\mathbb{G}_2$ be a Lipschitz map between Carnot groups. Then for almost every $x\in\mathbb{G}_1$, the sequence of maps*

$$\delta_\lambda\circ(L_{f(x)^{-1}}\circ f\circ L_x)\circ\delta_{\lambda^{-1}}$$

converges uniformly on compact sets, as $\lambda\to\infty$, to a Lie group homomorphism $Df(x)\colon\mathbb{G}_1\to\mathbb{G}_2$ that commutes with the dilations.

The following is an immediate consequence of Pansu's theorem, in the case that $\mathbb{G}_2=\mathbb{R}$. Let $n=\dim(V_1)$ be the vector space dimension of the horizontal layer of $\mathfrak{g}$. There is a natural "horizontal projection"

$$\pi\colon\mathbb{G}\to V_1\simeq\mathbb{R}^n$$

obtained by composing $\exp^{-1}$ with the vector space projection $P\colon\mathfrak{g}\to V_1$. If $f\colon\mathbb{G}\to\mathbb{R}$ is a Lipschitz function, then for almost every $x\in\mathbb{G}$, the mapping $Df(x)$ factors as $Df(x)=A\circ\pi$, where $A\colon V_1\to\mathbb{R}$ is linear.

The following lemma summarizes the additional basic properties of π that we will need below.

Lemma 2.8 *Let $\mathbb{G}$ be a Carnot group whose horizontal layer V_1 has dimension n, and let $\pi\colon\mathbb{G}\to V_1\simeq\mathbb{R}^n$ be the associated horizontal projection.*

(i) *$\mathbb{G}$ is a Lipschitz differentiability space with n–dimensional chart $(\mathbb{G},\pi)$.*

(ii) *π is a group homomorphism.*

(iii) *π commutes with dilations: $\pi(\delta_\lambda(x)) = \lambda\pi(x)$.*

(iv) *For all $x \in \mathbb{G}$, every element of $\mathrm{Tan}(\mathbb{G}, x, \pi)$ is isometric to $(\mathbb{G}, 0, \pi)$.*

(v) *π is a Lipschitz quotient map onto $V_1 \simeq \mathbb{R}^n$.*

(vi) *If $\gamma\colon [0,1] \to \mathbb{G}$ is a nonconstant Lipschitz curve, then $\pi \circ \gamma$ is nonconstant.*

(vii) *If $\gamma\colon [0,1] \to V_1$ is a Lipschitz curve and $x \in \pi^{-1}(\gamma(0))$, then there is a unique Lipschitz curve $\widetilde{\gamma}\colon [0,1] \to \mathbb{G}$ such that $\pi(\widetilde{\gamma}(t)) = \gamma(t)$ and $\widetilde{\gamma}(0) = x$.*

Proof That $\mathbb{G}$ is a Lipschitz differentiability space follows from the fact that all Carnot groups are PI spaces [29; 27] as well as Cheeger's theorem that all PI spaces are Lipschitz differentiability spaces [16]. That $\pi\colon \mathbb{G} \to \mathbb{R}^n$ serves as a global chart can be deduced from Pansu's theorem and the fact that for every group homomorphism $L\colon \mathbb{G} \to \mathbb{R}$ that commutes with dilations, there is a linear map $A\colon V_1 \to \mathbb{R}$ such that $L = A \circ \pi$. Here we use that π is, in fact, a Lipschitz map.

The Baker–Campbell–Hausdorff formula shows that π is a group homomorphism. Indeed,

$$\pi(xy) = P(\exp^{-1}(xy)) = P(\exp^{-1}(x) + \exp^{-1}(y) + v),$$

where $v \in [\mathfrak{g}, \mathfrak{g}] = V_2 \oplus \cdots \oplus V_s$. Thus,

$$\pi(xy) = P(\exp^{-1}(x)) + P(\exp^{-1}(y)) = \pi(x) + \pi(y).$$

That π commutes with the dilations on $\mathbb{G}$ and V_1 follows directly from the fact that dilations act on V_1 by simple scalings, as in (2-2). Part (iv) then follows immediately from parts (ii) and (iii).

Part (v) can be seen as follows. First of all, by Remark 2.6, it suffices to assume that $d = d_{\mathrm{CC}}$. By translation and dilation invariance, it further suffices to show that $\pi(B(0,1))$ contains an open ball around 0 in $V_1 \simeq \mathbb{R}^n$. Each element $v_1 \in B_{V_1}(0,1) \subseteq V_1$ gives rise to an element $x \in B(0,1)$ by exponentiating $v_1 \oplus 0 \oplus \cdots \oplus 0$. Since $\pi(x) = v_1$, we see that $\pi(B(0,1)) \supseteq B_{V_1}(0,1)$.

Part (vi) is immediate from the definition of the Carnot–Carathéodory metric and Remark 2.6.

Finally, property (vii) follows from [46, Proposition 2.3], using that γ is absolutely continuous. To show that the lift $\widetilde{\gamma}$ is Lipschitz, and not just absolutely continuous, one should use the fact that the horizontal derivative of $\widetilde{\gamma}(t)$ coincides almost everywhere with the derivative of $\gamma(t)$, and so is essentially bounded. □

We saw in Corollary 2.4 that tangents of Lipschitz differentiability spaces embedded in Euclidean space are themselves just Euclidean subspaces. For the appropriate generalization to Lipschitz differentiability spaces embedded in Carnot groups, we need the correct notion of a Carnot subgroup. This is given by the following definition:

Definition 2.9 Let $\mathbb{G}$ be a Carnot group with Lie algebra $\mathfrak{g}$ and horizontal layer $V_1 \subseteq \mathfrak{g}$. Let $V \subset V_1$ be a vector subspace, and let $\mathfrak{h} \subset \mathfrak{g}$ be the stratified Lie subalgebra generated by V. The homogeneous subgroup $H = \exp(\mathfrak{h}) \subset \mathbb{G}$ is called the *Carnot subgroup generated by* V.

Note that H is itself a Carnot group, and any homogeneous metric d on $\mathbb{G}$ restricts to a homogeneous metric on H. Moreover, H is rectifiably connected in this restricted metric.

3 Symmetric spaces and Gromov hyperbolic groups

Let X be a noncompact rank-one symmetric space, so that X is $\mathbb{H}^n_{\mathbb{R}}$, $\mathbb{H}^n_{\mathbb{C}}$, $\mathbb{H}^n_{\mathbb{Q}}$ or $\mathbb{H}^2_{\mathbb{O}}$ for some $n \geq 2$, where $\mathbb{Q}$ and $\mathbb{O}$ denote the quaternion and octonion division algebras. As a convention, we normalize the Riemannian metric on X to have maximal sectional curvature equal to -1. With the induced length metric d_X, the metric space (X, d_X) is therefore $\mathrm{CAT}(-1)$. We remind the reader that X is homogeneous and isotropic, so $\mathrm{Isom}(X)$ acts transitively on the unit tangent bundle of X.

For a discrete group Γ, an isometric and properly discontinuous action $\Gamma \curvearrowright X$ is said to be *convex-cocompact* if there is a convex, Γ–invariant subset $C(\Gamma) \subset X$ with $C(\Gamma)/\Gamma$ compact. This is equivalent to the seemingly weaker property that, for any point $p \in X$, the orbit map $g \mapsto g(p)$ gives a quasi-isometric embedding of Γ into X (see [9, Section 1.8; 24, Section 3]). As $C(\Gamma)$ is Gromov hyperbolic, the Švarc–Milnor lemma implies that Γ is necessarily a finitely generated hyperbolic group. From now on, we will always assume that Γ is nonelementary, ie is not finite and not virtually cyclic.

3.1 Visual boundaries of rank-one symmetric spaces

Let ∂X denote the visual boundary of X, which is a topological sphere of dimension $\dim_{\mathbb{R}}(X) - 1$. There are two natural classes of metrics on this boundary on which we focus. First, the visual metrics on ∂X are defined by

$$d_p(x, y) = e^{-(x,y)_p} \quad \text{for } x, y \in \partial X$$

for any $p \in X$, where $(x, y)_p$ is the Gromov product of x and y based at p

[9, Section 2.5]. Every element of $\mathrm{Isom}(X)$ extends to a homeomorphism of ∂X that is Möbius with respect to any visual metric.

The parabolic visual metrics are similar but are better suited to the parabolic models for X. Namely, for any $\omega \in \partial X$ and $q \in X$, define

$$d_{\omega,q}(x, y) = e^{-(x,y)_{\omega,q}} \quad \text{for } x, y \in \partial X \backslash \{\omega\},$$

where

$$(x, y)_{\omega,q} = \lim_{p \to \omega} ((x, y)_p - d_X(p, q))$$

is a limit taken along a geodesic ray in X that is asymptotic to ω. This is a metric on $\partial X \backslash \{\omega\}$, which is obtained as a limit of rescaled visual metrics d_p as p tends toward ω nontangentially. We refer the reader to [7, Section 2] for details, noting that the definition we give for $(x, y)_{\omega,q}$ is a consequence of [7, Lemma 2.1].

Each visual metric is Möbius equivalent to each parabolic visual metric on their common domains [7, Lemma 2.3]. An immediate consequence is that any two parabolic visual metrics are Möbius equivalent on their common domains. Moreover, for fixed $\omega \in \partial X$, the metrics $d_{\omega,q}$ and $d_{\omega,q'}$ differ by a scalar multiple for any $q, q' \in X$. This follows from the fact that $\lim_{p \to \omega}(d_X(p, q) - d_X(p, q'))$ exists. These two metrics coincide precisely when q and q' lie on the same horosphere based at ω.

We should note that the above discussion holds equally well for boundaries of $\mathrm{CAT}(-1)$ metric spaces. The important point for us is that, when X is noncompact rank-one symmetric, the boundary has much additional structure. Namely, given a point $\omega \in \partial X$, there is a natural identification of $\partial X \backslash \{\omega\}$ with a Carnot group $\mathbb{G}$. Here, the horizontal distribution on $\mathbb{G}$ arises from vectors that are tangent to the lines and circles in $\partial X \backslash \{\omega\}$ formed by isometric copies of $\mathbb{H}^2_{\mathbb{R}}$ in X. Moreover, the subgroup of $\mathrm{Isom}(X)$ that fixes ω corresponds to the collection of affine transformations of $\mathbb{G}$. In particular, this includes all left translations and dilations.

Lemma 3.1 *Each parabolic visual metric $d_{\omega,q}$ is a homogeneous distance on $\mathbb{G} = \partial X \backslash \{\omega\}$.*

Proof That $d_{\omega,q}$ induces the Euclidean topology is a direct consequence of the standard fact that any visual metric d_p induces the spherical topology on ∂X. The other two properties are consequences of the identity

$$d_{\omega,g(q)}(g(x), g(y)) = d_{\omega,q}(x, y) \quad \text{for } x, y \in \partial X \backslash \{\omega\}$$

whenever $g \in \mathrm{Isom}(X)$ fixes ω.

Indeed, every left translation of $\partial X\backslash\{\omega\}$ is the boundary map of some element $g \in \mathrm{Isom}(X)$ that fixes ω and preserves the horospheres in X that are based at ω. As $g(q)$ and q lie on the same horosphere, we have $d_{\omega,g(q)} = d_{\omega,q}$, which shows that $d_{\omega,q}$ is left-invariant.

Similarly, every dilation δ_λ of $\partial X\backslash\{\omega\}$ is the boundary map of some element $g \in \mathrm{Isom}(X)$ that acts as a translation by distance $\log\lambda$ along a geodesic in X that is asymptotic to ω. This means that

$$\lim_{p\to\omega}\big(d_X(p,q) - d_X(p,g(q))\big) = \log\lambda,$$

so we obtain $d_{\omega,g(q)} = \lambda d_{\omega,q}$. This shows that $d_{\omega,q}$ is 1–homogeneous. □

Remark 3.2 Identifying $\partial X\backslash\{\omega\}$ with a horosphere in X based at ω, one can obtain a sub-Riemannian Carnot–Carathéodory metric on $\partial X\backslash\{\omega\}$ as a limit of Riemannian metrics (see [14]). In general, the parabolic visual metrics $d_{\omega,q}$ are not geodesic, and so will not coincide with the sub-Riemannian metric. However, the parabolic visual metrics are better suited for our work in Section 5.

We should note that the more general identity

$$d_{g(\omega),g(q)}(g(x),g(y)) = d_{\omega,q}(x,y) \quad \text{for } x,y \in \partial X\backslash\{\omega\}$$

holds for any $g \in \mathrm{Isom}(X)$. Using that $\mathrm{Isom}(X)$ acts transitively on the unit tangent bundle of X, for any two pairs $\omega,\omega' \in \partial X$ and $q,q' \in X$, there is $g \in \mathrm{Isom}(X)$ with $g(\omega) = \omega'$ and $g(q) = q'$. Thus, the parabolic boundaries $(\partial X\backslash\{\omega\}, d_{\omega,q})$ and $(\partial X\backslash\{\omega'\}, d_{\omega',q'})$ are isometrically equivalent.

We therefore consider the Carnot group $\mathbb{G}$, equipped with any parabolic metric $d = d_{\omega,q}$, to be a model for the local geometry of $(\partial X, d_p)$. By our discussion above, it is clear that $(\mathbb{G}, d)$ is Möbius equivalent to $(\partial X\backslash\{\omega\}, d_p)$, regardless of the choice of $\omega \in \partial X$. In particular, this means that $(\partial X, d_p)$ is locally bi-Lipschitz equivalent to $(\mathbb{G}, d)$, so it is also locally bi-Lipschitz equivalent to $\mathbb{G}$ equipped with any homogeneous distance.

Of course, the Carnot groups that locally model boundaries of noncompact rank-one symmetric spaces are a special sort. They are either Euclidean, Heisenberg, or of "Heisenberg type":

(i) If $X = \mathbb{H}^n_{\mathbb{R}}$, then $\mathbb{G} = \mathbb{R}^{n-1}$ is Euclidean space.

(ii) If $X = \mathbb{H}^n_{\mathbb{C}}$, then $\mathbb{G} = \mathcal{H}^{n-1}_{\mathbb{C}}$ is the n^{th} Heisenberg group.

(iii) If $X = \mathbb{H}^n_{\mathbb{Q}}$, then $\mathbb{G} = \mathcal{H}^{n-1}_{\mathbb{Q}}$ is the n^{th} quaternionic Heisenberg group.

(iv) If $X = \mathbb{H}^2_{\mathbb{O}}$, then $\mathbb{G} = \mathcal{H}^1_{\mathbb{O}}$ is the first octonionic Heisenberg group.

Together, these Carnot groups form the class of *Iwasawa groups*. In each case, isometries of X act on ∂X by conformal maps, ie smooth maps for which the restriction of the derivative to the horizontal layer is a similarity. For $\mathbb{H}^n_{\mathbb{R}}$ this is classical, and the boundary action is by (classical) Möbius transformations; for $\mathbb{H}^n_{\mathbb{C}}$, this is shown in [37, page 328]; the other cases follow from [43, Corollary 11.2; 15, Corollary 7.2].

If $\Gamma \curvearrowright X$ is a convex-cocompact action with Γ nonelementary, the limit set $\Lambda(\Gamma) \subset \partial X$ is defined to be the visual boundary of any convex, Γ–invariant subset $C(\Gamma) \subset X$ for which $C(\Gamma)/\Gamma$ is compact. Equivalently, if $p \in X$ is any point, $\Lambda(\Gamma)$ is the image of the boundary map $\partial\Gamma \to \partial X$ induced by the quasi-isometric orbit embedding $\Gamma \to X$. In particular, $\Lambda(\Gamma)$ is a closed subset, and the Möbius action $\Gamma \curvearrowright \partial X$ leaves $\Lambda(\Gamma)$ invariant. Thus, we obtain a natural Möbius action $\Gamma \curvearrowright \Lambda(\Gamma)$, when we equip $\Lambda(\Gamma)$ with the restriction of any visual metric on ∂X. It is known that the corresponding Hausdorff measure is Ahlfors regular [9, Section 2.7].

3.2 Boundaries of Gromov hyperbolic groups

Much of the content of this subsection is background and will not be needed in the remainder of the paper. However, we believe it should be recorded in the literature, and some of it is needed in the proof of our main result. For further background and terminology about Gromov hyperbolic groups, we refer the reader to [8].

Let Γ be a nonelementary Gromov hyperbolic group, by which we mean in particular that Γ is finitely generated. The visual boundary $\partial\Gamma$ is perfect and compact, and it admits a collection of visual Gromov metrics, each of which is Ahlfors regular [21]. Any two such metrics are quasisymmetrically equivalent. Moreover, the action of Γ on itself by left multiplication extends naturally to a boundary action $\Gamma \curvearrowright \partial\Gamma$ that is uniformly quasi-Möbius (with linear distortion function $\eta(t)$) with respect to any visual Gromov metric [35, Section 6].

We will use $\mathcal{J}_{\mathrm{AR}}(\partial\Gamma)$ to denote the Ahlfors regular conformal gauge of $\partial\Gamma$ that contains these metrics, ie the collection of all Ahlfors regular metric spaces quasisymmetric to $\partial\Gamma$. By definition,

$$\operatorname{confdim}_{\mathrm{AR}}(\partial\Gamma) = \inf\{\operatorname{Hdim}(Z) : Z \in \mathcal{J}_{\mathrm{AR}}(\partial\Gamma)\}.$$

Many quasi-isometric uniformization statements about Γ boil down to finding a highly regular metric in $\mathcal{J}_{\mathrm{AR}}(\partial\Gamma)$. The following gives a list of equivalent notions for "highly regular". In what follows, all of the metric spaces that appear will be Ahlfors regular, and we endow them with the corresponding Hausdorff measure.

The following result is surely known to experts, but we include it as a useful summation. The definitions of Poincaré inequalities and Loewner spaces can be found in [27] or [26].

Theorem 3.3 *For $Z \in \mathcal{J}_{\mathrm{AR}}(\partial\Gamma)$ of Hausdorff dimension $Q > 1$, the following are equivalent:*

(i) *Z admits a $(1, p)$–Poincaré inequality for some $p \geq 1$.*

(ii) *Z is a Lipschitz differentiability space.*

(iii) *Z has Ahlfors regular conformal dimension equal to Q.*

(iv) *Z is a Q–Loewner space.*

(v) *Z admits a path family of positive p–modulus for some $p \geq 1$.*

Proof As Z is in $\mathcal{J}_{\mathrm{AR}}(\partial\Gamma)$ and has Hausdorff dimension Q, it is Ahlfors Q–regular. We first show that properties (i) through (iv) are equivalent.

The implication (i) implies (ii) is a consequence of the main theorem of Cheeger [16], using the fact that Z is Ahlfors regular and hence a doubling metric measure space.

To see that (ii) implies (iii), we first note that [18, Theorem 1.15] shows that Z has a tangent $\widehat{Z}$ that admits an *Alberti representation* supported on geodesic lines. (Here, $\widehat{Z}$ is also Ahlfors Q–regular, so we use the corresponding Hausdorff measure.) We refer the reader to [18] for the definition of an Alberti representation; all we will need is the fact that the geodesic lines in the support of this Alberti representation constitute a path family in $\widehat{Z}$ that has positive 1–modulus. Restricting these geodesics to their intersections with a large fixed ball in $\widehat{Z}$ gives a path family of positive 1–modulus inside of a compact set. By Hölder's inequality, this path family has positive Q–modulus as well. It now follows from [40, Proposition 4.1.8] and the fact that Z itself is Ahlfors Q–regular, that Z has Ahlfors regular conformal dimension equal to Q.

That (iii) implies (iv) is a consequence of [8, Theorem 1.3] (and the remark following the statement of that theorem).

Since Z is Ahlfors Q–regular, property (iv), that Z is Q–Loewner, implies that Z has a $(1, Q)$–Poincaré inequality and hence property (i) [27, Theorem 5.12].

It remains to show that (v) is equivalent to the other properties. Of course, property (iv) implies the existence of a path of positive Q–modulus and hence property (v). On the other hand, if Z admits a path family with positive p–modulus for some $p \geq 1$, then [33, Theorem 4.0.5] shows that Z has a tangent with a path family of positive 1–modulus, hence of positive Q–modulus as above. It then follows from [40, Corollary 6.1.8] that Z has Ahlfors regular conformal dimension equal to Q, ie that (iii) holds. □

Remark 3.4 When $\Gamma \curvearrowright X$ is a convex-cocompact action on a noncompact rank-one symmetric space X, as in the previous subsection, the boundary homeomorphism between $\partial\Gamma$ and $\Lambda(\Gamma)$ is quasisymmetric. As $\Lambda(\Gamma)$ is Ahlfors regular, we have $\Lambda(\Gamma) \in \mathcal{J}_{\mathrm{AR}}(\partial\Gamma)$. Consequently,

$$\mathrm{Hdim}(\Lambda(\Gamma)) \geq \mathrm{confdim}_{\mathrm{AR}}(\Lambda(\Gamma)) = \mathrm{confdim}_{\mathrm{AR}}(\partial\Gamma),$$

which is the inequality that appears in Theorems 1.3 and 1.4.

The case of equality, $\mathrm{Hdim}(\Lambda(\Gamma)) = \mathrm{confdim}_{\mathrm{AR}}(\partial\Gamma)$, means precisely that condition (iii) in Theorem 3.3 holds for $Z = \Lambda(\Gamma)$ and $Q = \mathrm{confdim}_{\mathrm{AR}}(\partial\Gamma)$. Thus, if $Q > 1$, then the other conditions hold for $\Lambda(\Gamma)$ as well.

On the other hand, if $Q \leq 1$, there is only one possibility for the action $\Gamma \curvearrowright X$. Indeed, since $Q \leq 1$, the topological dimension of $\Lambda(\Gamma)$ is either 0 or 1. If the former, then $\partial\Gamma$ also has topological dimension 0, and Γ is virtually a free group by [31, Theorem 8.1]. This means that $\partial\Gamma$ is a uniformly perfect Cantor set, which is known to have Ahlfors regular conformal dimension equal to 0. Hence, $Q = 0$, which implies that $\Lambda(\Gamma)$ is finite, and so Γ is elementary, a contradiction. Thus, the topological dimension of $\Lambda(\Gamma)$ must be equal to 1, which then means that $Q = 1$ as well. Applying [7, Theorem 1.1] shows that Γ is virtually Fuchsian and the action $\Gamma \curvearrowright X$ stabilizes an isometric copy of $\mathbb{H}^2_{\mathbb{R}}$ in X.

These arguments justify our subsequent restriction to the case $Q > 1$ and our phrasing of Theorem 1.3 above (which, in [8], is stated only for $Q > 1$).

In the remainder of the section, we make some general remarks about the types of Lipschitz differentiability structures that can appear on boundaries of hyperbolic groups.

Let Γ be a Gromov hyperbolic group with $\operatorname{confdim}_{\mathrm{AR}}(\partial\Gamma) = Q > 1$. If there is $Z \in \mathcal{J}_{\mathrm{AR}}(\partial\Gamma)$ that is a Lipschitz differentiability space, then Theorem 3.3 guarantees that $\dim_H(Z) = Q$. In other words, the Hausdorff dimension of any highly regular metric on $\partial\Gamma$ depends only on the quasi-isometry class of Γ. Being slightly imprecise, one could simply refer to Q as the Hausdorff dimension of $\partial\Gamma$.

We claim that a similar statement holds for the dimension of the differentiability structure. In fact, we can establish something stronger.

Lemma 3.5 (i) *If $Z \in \mathcal{J}_{\mathrm{AR}}(\partial\Gamma)$ is a Lipschitz differentiability space, then the action $\Gamma \curvearrowright Z$, obtained by conjugating the boundary action $\Gamma \curvearrowright \partial\Gamma$ by a quasisymmetric homeomorphism between $\partial\Gamma$ and Z, is ergodic with respect to Q–dimensional Hausdorff measure.*

(ii) *If $Z_1, Z_2 \in \mathcal{J}_{\mathrm{AR}}(\partial\Gamma)$ are Lipschitz differentiability spaces, then for any Borel sets $U_1 \subset Z_1$ and $U_2 \subset Z_2$ of positive measure, there are positive-measure subsets $A_1 \subset U_1$ and $A_2 \subset U_2$ that are bi-Lipschitz equivalent.*

Proof To verify (i), first note that Z is Ahlfors Q–regular and Q–Loewner by Theorem 3.3. Let μ denote the Q–dimensional Hausdorff measure, which we may normalize to have $\mu(Z) = 1$.

The action $\Gamma \curvearrowright Z$ is uniformly quasi-Möbius and, hence, uniformly quasiconformal: there is $K < \infty$ for which each $g \in \Gamma$ acts as a K–quasiconformal homeomorphism of Z. Using that Z is Ahlfors Q–regular and Q–Loewner, this means that each $g \in \Gamma$ is absolutely continuous in measure [28, Corollary 8.15]. Hence, $\Gamma \curvearrowright Z$ is a measure-class-preserving action. Ergodicity of such an action means, as usual, that any Γ–invariant Borel set has measure 0 or 1.

To show this, we note that each $g \in \Gamma$ has a uniform density property (in the sense of [38, Section 6]) with uniform distortion control. Actually, we will use a slightly different property, but which is of the same spirit: there is a homeomorphism $\phi\colon [0,\infty) \to [0,\infty)$ such that

$$\mu(g(E \cap B)) \le \phi\left(\frac{\mu(E \cap B)}{\mu(B)}\right)$$

for every $g \in \Gamma$, every Borel set $E \subset Z$eand every ball $B \subset Z$. A minor modification of the proof of [38, Theorem 6.3] shows this easily, once we note that Z satisfies a $(1, p)$–Poincaré inequality for some $1 < p < Q$ by the main result of [34].

Now suppose that $E \subset Z$ is a Γ–invariant Borel set with $\mu(E) < 1$. Let $\epsilon > 0$ be arbitrary, and let $B \subset Z$ be a ball, centered at a point of density for $Z \backslash E$, with radius small enough that

$$\phi\left(\frac{\mu(E \cap B)}{\mu(B)}\right) < \frac{\epsilon}{2}.$$

By [6, Lemma 5.1], there are elements $g \in \Gamma$ for which $\operatorname{diam}(Z \backslash g(B))$ is arbitrarily small. In particular, we can find $g \in \Gamma$ for which

$$\mu(Z \backslash g(B)) < \frac{\epsilon}{2}.$$

Using that E is Γ–invariant, we have $E = g(E) \subset g(E \cap B) \cup (Z \backslash g(B))$, and thus

$$\mu(E) \leq \mu(g(E \cap B)) + \mu(Z \backslash g(B)) < \phi\left(\frac{\mu(E \cap B)}{\mu(B)}\right) + \frac{\epsilon}{2} < \epsilon.$$

As $\epsilon > 0$ was arbitrary, we see that $\mu(E) = 0$. Hence, $\Gamma \curvearrowright Z$ is ergodic.

Let us now verify part (ii). Again, we know that Z_1 and Z_2 are both Ahlfors Q–regular and Q–Loewner. Moreover, there is a quasisymmetric homeomorphism $f\colon Z_1 \to Z_2$, which is absolutely continuous in measure by [28, Corollary 8.15]. In particular, the image $f(U_1)$ has positive measure in Z_2. By part (i), the action $\Gamma \curvearrowright Z_2$ is ergodic, so there is $g \in \Gamma$ for which $g(f(U_1)) \cap U_2$ has positive measure. Replacing f by the composition $g \circ f$, we may suppose without loss of generality that $f(U_1) \cap U_2$ has positive measure. This also means that $U_1 \cap f^{-1}(U_2)$ has positive measure.

It is shown in [28, Section 10] that if $f\colon Z_1 \to Z_2$ is a quasisymmetric homeomorphism between Ahlfors Q–regular, Q–Loewner spaces, then Z_1 can be covered, up to a set of measure zero, by Borel sets on which f is Lipschitz. In particular, given any positive-measure subset $B \subset Z_1$, there is a positive-measure set $A \subset B$ on which f is Lipschitz. Using the same argument for f^{-1}, along with the fact that f and f^{-1} preserve sets of measure zero, it is easy to see that we can actually find a positive-measure subset $A \subset B$ on which f is bi-Lipschitz. Applying this fact to the set $U_1 \cap f^{-1}(U_2) \subset Z_1$, we find a positive-measure subset $A_1 \subset U_1 \cap f^{-1}(U_2)$ on which f is bi-Lipschitz. Then set $A_2 = f(A_1) \subset U_2$. □

An immediate consequence of Lemma 3.5(ii) is that there is a unique integer k for which each chart in any Lipschitz differentiability space $Z \in \mathcal{J}_{\mathrm{AR}}(\partial\Gamma)$ has dimension k. It makes sense to call this integer the analytic dimension of $\partial\Gamma$, and once again it depends only on the quasi-isometry class of Γ. Thus, in the case that $\mathcal{J}_{\mathrm{AR}}(\partial\Gamma)$ contains a Lipschitz differentiability space, there are three natural well-defined dimensions to

consider: the topological dimension of $\partial\Gamma$, the Hausdorff dimension of $\partial\Gamma$ and the analytic dimension of $\partial\Gamma$.

At the same time, it should be quite rare for $\mathcal{J}_{\mathrm{AR}}(\partial\Gamma)$ to contain a Lipschitz differentiability space. Currently, the only known examples come from uniform lattices in noncompact rank-one symmetric spaces and uniform lattices in certain types of hyperbolic buildings [11]. It makes more sense to look for rigidity phenomena for such groups Γ. One might expect the topological, Hausdorff and analytic dimensions of $\partial\Gamma$ to function as characteristic quantities for rigidity.

There are some intimations toward this type of rigidity when Γ is a manifold group. If $\Gamma = \pi_1(M)$ with M a closed, negatively curved Riemannian manifold, of dimension at least 3, then the action by deck transformations $\Gamma \curvearrowright \widetilde{M}$ on the universal Riemannian cover has compact quotient. Rescaling the metric, we may assume that the maximal sectional curvature of $\widetilde{M}$ equals -1. The visual boundary $\partial\widetilde{M}$, equipped with any visual metric, is then in $\mathcal{J}_{\mathrm{AR}}(\partial\Gamma)$.

If $\partial\widetilde{M}$ has a path family of positive p–modulus for some $p \geq 1$, then a theorem of Connell shows that $\widetilde{M}$ is isometric to a noncompact rank-one symmetric space (Theorem 4.3 of [20] for Patterson–Sullivan measures). This symmetric space is uniquely determined by the topological and Hausdorff dimensions of $\partial\widetilde{M}$, as it is locally modeled by an Iwasawa group. It is not clear, however, whether the same conclusion holds if one assumes only that $\mathcal{J}_{\mathrm{AR}}(\partial\Gamma)$ contains a Lipschitz differentiability space.

Finally, let us record an analog of the Patterson–Sullivan case in Connell's result that uses some of the ideas we will see below.

Proposition 3.6 *Let $\Gamma \curvearrowright X$ be a convex-cocompact action on a* $\mathrm{CAT}(-1)$ *space X with* $\mathrm{Hdim}(\Lambda(\Gamma)) > 1$. *Suppose that $\Lambda(\Gamma)$ admits a path family of positive p–modulus for some $p \geq 1$. Then every geodesic in X whose endpoints lie in $\Lambda(\Gamma)$ is contained in an isometrically embedded copy of $\mathbb{H}^2_{\mathbb{R}}$ in X for which the boundary circle lies entirely in $\Lambda(\Gamma)$.*

In the setting of Connell's theorem, the action of Γ on $X = \widetilde{M}$ is cocompact, so the conclusion of Proposition 3.6 holds for all geodesics in $\widetilde{M}$. This means that $\widetilde{M}$ has hyperbolic rank at least 1, and a powerful theorem of Hamenstädt then implies that $\widetilde{M}$ is symmetric [25].

Proof By a theorem of Bourdon [10, Theorem 0.1], it suffices to show that any two points in $\Lambda(\Gamma)$ can be joined by a Möbius circle that lies in $\Lambda(\Gamma)$. As Möbius

circles in ∂X are closed under nontrivial limits, it actually suffices to show that for a dense set of points $\omega \in \Lambda(\Gamma)$, for every $\eta \in \Lambda(\Gamma)\backslash\{\omega\}$, there is a Möbius circle in $\Lambda(\Gamma)$ containing ω and η. Recall, though, that the boundary action $\Gamma \curvearrowright \Lambda(\Gamma)$ is Möbius, and the orbit of every point is dense. Thus, it suffices to prove the previous statement for a single point $\omega \in \Lambda(\Gamma)$. A natural rephrasing of this statement is that every $\eta \in \Lambda(\Gamma)\backslash\{\omega\}$ is contained in an isometric copy of $\mathbb{R}$ in the parabolic limit set $(\Lambda(\Gamma)\backslash\{\omega\}, d_{\omega,q})$ for some $q \in X$. Let us prove this formulation.

By Theorem 3.3, we know that $\Lambda(\Gamma)$ is a Lipschitz differentiability space. A result of Cheeger, Kleiner and Schioppa, Proposition 2.3 above, guarantees that $\Lambda(\Gamma)$ has a tangent Y for which there is a metric submersion $Y \to (\mathbb{R}^k, ||\cdot||)$ onto a normed space. A theorem of Bonk and Kleiner [7, Proposition 3.1] shows that there are points $\omega \in \Lambda(\Gamma)$ and $q \in X$ for which Y is isometrically equivalent to $(\Lambda(\Gamma)\backslash\{\omega\}, d_{\omega,q})$. Thus, there is a metric submersion

$$f\colon (\Lambda(\Gamma)\backslash\{\omega\}, d_{\omega,q}) \to (\mathbb{R}^k, ||\cdot||).$$

For $\eta \in \Lambda(\Gamma)\backslash\{\omega\}$, let ℓ be a geodesic line in $\mathbb{R}^k$ through $f(\eta)$, eg the line in the first coordinate direction. By Lemma 4.3 below, there is a lift of ℓ to a geodesic line in $\Lambda(\Gamma)\backslash\{\omega\}$ that contains η, as desired. □

It is desirable to understand better the global geometry of X that can arise in this setting, even in the case that $\Gamma \curvearrowright X$ is cocompact. One can consider Theorem 1.4 to be a description of what happens when the ambient space X is not only CAT(-1) but is in fact symmetric.

4 Tangents of Lipschitz differentiability spaces in Carnot groups

The main result of this section is Theorem 4.1, which essentially shows that if a Lipschitz differentiability space lies inside a Carnot group, then its tangents are Carnot subgroups. Theorem 4.1 is an important piece in the proof of Theorem 1.4, but also has other interesting nonembedding consequences for Lipschitz differentiability spaces, as explained in Section 4.1.

Let $\mathbb{G}$ be a Carnot group, and let d be a homogeneous distance on $\mathbb{G}$. Let $X \subset \mathbb{G}$ be a subset for which there is a Radon measure μ on X for which (X, d, μ) is a Lipschitz differentiability space. Recall the definition of Carnot subgroups given in

Definition 2.9. Note that X is automatically metrically doubling, as a subset of the metrically doubling space $\mathbb{G}$.

Theorem 4.1 *Let (U,ϕ) be a k–dimensional chart for X, and assume that for almost every $x\in U$, each $(\widehat{X},\widehat{x})\in \operatorname{Tan}(X,x)$ is rectifiably connected. Then for almost every $x\in U$, the set $\operatorname{Tan}(X,x)$ consists of one element, which is canonically isometric to a Carnot subgroup of $\mathbb{G}$ generated by a fixed k–dimensional subspace of the horizontal layer of $\mathbb{G}$.*

Remark 4.2 If (X,d,μ) is a PI space (ie is doubling and satisfies a Poincaré inequality in the sense of [27]), then the assumption that each $(\widehat{X},\widehat{x})\in \operatorname{Tan}(X,x)$ is rectifiably connected can be omitted in Theorem 4.1. This is because each element of $\operatorname{Tan}(X,x)$ will be quasiconvex and hence rectifiably connected. Thus, Theorem 1.5 from the introduction, which is the only case of Theorem 4.1 needed for the proof of Theorem 1.4, follows immediately.

The assumption that each tangent is rectifiably connected in Theorem 4.1 can also be omitted if (X,d,μ) is an RNP Lipschitz differentiability space, in the sense of [3]. This follows from Corollary 9.4 of [3].

The following path lifting lemma for Lipschitz quotients is taken from [5, Lemma 4.4] (equivalently [30, Lemma 2.2]). Although stated there only for Euclidean domains, the proof works equally well in the context below. Recall that a metric space is *proper* if each closed ball in the space is compact. Every complete, doubling metric space is proper.

Lemma 4.3 *Let X be a proper metric space and Y a metric space. Let $F\colon X\to Y$ be a Lipschitz quotient with co-Lipschitz constant $c>0$, and let $\gamma\colon[0,T]\to Y$ be a 1–Lipschitz curve with $\gamma(0)=F(x)$. Then there is a $(1/c)$–Lipschitz curve $\widetilde{\gamma}\colon[0,T]\to X$ such that $\widetilde{\gamma}(0)=x$ and $F\circ\widetilde{\gamma}=\gamma$.*

We now let $\pi\colon\mathbb{G}\to V_1\simeq\mathbb{R}^n$ be the global differentiability chart for $\mathbb{G}$, and let $\pi=(\pi_1,\dots,\pi_n)$ denote its coordinates. Recall that $(U,\phi\colon X\to\mathbb{R}^k)$ is a differentiability chart for the space (X,d,μ) contained in $\mathbb{G}$.

Lemma 4.4 *The set U can be covered by a finite number of charts, with chart maps of the form*

$$(\pi_{i_1},\dots,\pi_{i_k})\colon X\to\mathbb{R}^k$$

for some choice of k distinct indices $i_1,\dots,i_k$.

Proof Let $\iota\colon X \hookrightarrow \mathbb{G}$ be the inclusion map, which is an isometry. Postcomposing with $\pi\colon \mathbb{G} \to \mathbb{R}^n$ gives a Lipschitz map from X to $\mathbb{R}^n$, which can be differentiated with respect to the chart $\phi\colon X \to \mathbb{R}^k$ for U. Thus, for almost every $x \in U$, there is a unique linear map

$$D\iota(x)\colon \mathbb{R}^k \to \mathbb{R}^n$$

for which

$$\pi(y) - \pi(x) = D\iota(x)(\phi(y) - \phi(x)) + o(d(x,y)) \quad \text{for } y \in X. \tag{4-1}$$

For almost every $x \in U$, if $(\widehat{X}, \hat{x}) \in \mathrm{Tan}(X, x)$, then the corresponding tangent maps $\hat{\phi}\colon \widehat{X} \to \mathbb{R}^k$ and $\hat{\iota}\colon \widehat{X} \to \mathbb{G}$ satisfy

$$\pi \circ \hat{\iota} = D\iota(x) \circ \hat{\phi}.$$

Here we consider $\widehat{X}$ as a subset of $\mathbb{G}$, so that $\hat{\iota}$ is just the inclusion map again. Moreover, $\hat{\phi}\colon \widehat{X} \to \mathbb{R}^k$ is a Lipschitz quotient mapping.

Consider any nonzero vector $v \in \mathbb{R}^k$. By Lemma 4.3, there is a Lipschitz curve $\ell\colon [0,1] \to \widehat{X} \subset \mathbb{G}$ with

$$\hat{\phi} \circ \ell(t) = tv.$$

In particular, we have

$$\pi \circ \ell(t) = D\iota(x)(tv) = t \cdot D\iota(x)(v).$$

Since ℓ is a nonconstant curve in $\mathbb{G}$, it must be that $\pi \circ \ell$ is nonconstant (Lemma 2.8(vi)). Hence, $D\iota(x)(v) \neq 0$. As v was arbitrary, we conclude that $D\iota(x)$ is injective, and in particular that $k \leq n$.

Given $i_1 < \cdots < i_k \in \{1, \ldots, n\}$, let $U_{i_1,\ldots,i_k}$ denote the subset of U on which the $k \times k$ minor of $(D\iota)(x)$ defined by the coordinates $i_1, \ldots, i_k$ is invertible. Note that almost every $x \in U$ is in some such set, since $(D\iota)(x)$ is injective for almost every $x \in U$. For $x \in U_{i_1,\ldots,i_k}$, let $A(x)\colon \mathbb{R}^n \to \mathbb{R}^k$ be the unique linear mapping with

$$\ker(A) = \mathrm{span}(\{e_1, \ldots, e_n\} \setminus \{e_{i_1}, \ldots, e_{i_k}\}) \tag{4-2}$$

such that

$$A(x) \cdot (D\iota)(x) = \mathrm{Id}_k. \tag{4-3}$$

Applying (4-3) to (4-1), we see that

$$\phi(y) - \phi(x) = A(x) \cdot (\pi(y) - \pi(x)) + o(d(x,y)) \quad \text{for } y \in X. \tag{4-4}$$

From (4-4) and (4-2), it follows that a Lipschitz function $f\colon X \to \mathbb{R}$ that is differentiable with respect to ϕ at $x \in U_{i_1,\dots,i_k}$, with unique derivative, is differentiable with respect to the restrictions $(\pi_{i_1}|_X, \dots, \pi_{i_k}|_X)$ with unique derivative.

Hence, these form a k–dimensional chart map for $U_{i_1,\dots,i_k}$. □

We are now ready to prove Theorem 4.1.

Proof of Theorem 4.1 Using the previous lemma, and passing to subsets, we may assume without loss of generality that U has chart map

$$(\pi_1, \dots, \pi_k)\colon X \to \mathbb{R}^k \subset \mathbb{R}^n.$$

For almost every $x \in U$, we know that the other coordinates $\pi_{k+1}, \dots, \pi_n$ are linear combinations of $\pi_1, \dots, \pi_k$, up to first order on X near x. Fix such a point x, and assume also that every $(\widehat{X}, \hat{x}) \in \operatorname{Tan}(X, x)$ is rectifiably connected and all tangent maps $(\hat{\pi}_1, \dots, \hat{\pi}_k)\colon \widehat{X} \to \mathbb{R}^k$ are Lipschitz quotients. By assumption, such points form a set of full measure in U.

Fix $(\widehat{X}, \hat{x}) \in \operatorname{Tan}(X, x)$. We canonically identify $\widehat{X}$ with a closed subset of $\mathbb{G}$ that is a limit of rescalings of $L_{-x}(X)$, with $\hat{x} = 0$ and $\hat{\pi}_i = \pi_i$. Moreover, as the coordinates $\pi_{k+1}, \dots, \pi_n$ were linear combinations of $\pi_1, \dots, \pi_k$, up to first order on X near x, we see that $\pi_{k+1}|_{\widehat{X}}, \dots, \pi_n|_{\widehat{X}}$ are precisely linear combinations of $\pi_1|_{\widehat{X}}, \dots, \pi_k|_{\widehat{X}}$ on $\widehat{X}$. In particular, there is an injective linear transformation $A\colon \mathbb{R}^k \to \mathbb{R}^n$ for which

$$p := \pi|_{\widehat{X}} = A \circ (\pi_1, \dots, \pi_k).$$

Let $V = p(\widehat{X}) = A(\mathbb{R}^k)$ be the corresponding linear subspace of $\mathbb{R}^n \simeq V_1$. It is a k–dimensional subspace because $(\pi_1, \dots, \pi_k)\colon \widehat{X} \to \mathbb{R}^k$ is surjective by the Lipschitz quotient property and A is injective. Furthermore, the map $p\colon \widehat{X} \to V$ is a Lipschitz quotient.

We claim that $\widehat{X}$ is the Carnot subgroup of $\mathbb{G}$ generated by $V \subset \mathbb{R}^n$. First, we show that it is a subgroup. Let $y, z \in \widehat{X}$, and let γ be a rectifiable curve in $\widehat{X}$ from 0 to z, so that $p \circ \gamma$ is a rectifiable curve in V from 0 to $p(z)$. Consider the rectifiable curve $L_y \circ \gamma$, which joins y to yz in $\mathbb{G}$. As

$$\pi \circ L_y \circ \gamma = L_{\pi(y)} \circ \pi \circ \gamma = L_{p(y)} \circ p \circ \gamma,$$

we see that $L_y \circ \gamma$ is the unique horizontal lift through π of the curve $L_{p(y)} \circ p \circ \gamma$ to $\mathbb{G}$, starting at y (recall the uniqueness in Lemma 2.8(vii)). Note that $L_{p(y)} \circ p \circ \gamma$

is a rectifiable curve in V, as $L_{p(y)}$ preserves this subspace. Applying Lemma 4.3 to the Lipschitz quotient $p\colon \widehat{X} \to V$, we obtain a lift of $L_{p(y)} \circ p \circ \gamma$ to a rectifiable curve in $\widehat{X}$ beginning at y. This lift through p is also a lift through π, and by uniqueness of horizontal lifts to $\mathbb{G}$, the two lifts must coincide. Hence, $L_y \circ \gamma$ is contained in $\widehat{X}$; in particular, its endpoint yz is in $\widehat{X}$.

Similar arguments show that $\widehat{X}$ is closed under inversion and dilation. If $y \in \widehat{X}$, then let γ be a rectifiable curve in $\widehat{X}$ from 0 to y. Note that $L_{y^{-1}} \circ \gamma$ is also rectifiable in $\mathbb{G}$. The reversal of this curve is therefore the unique lift of the reversal of $-\pi(y) + \pi \circ \gamma$ to $\mathbb{G}$ that begins at 0. As $-\pi(y) + \pi \circ \gamma$ is a rectifiable curve in V, the Lipschitz quotient property of p ensures that this lift lies in $\widehat{X}$. In particular, $y^{-1} \in \widehat{X}$. Finally, if δ_λ is a dilation for $\mathbb{G}$, then $\delta_\lambda \circ \gamma$ is again a horizontal curve in $\mathbb{G}$. As

$$\pi \circ \delta_\lambda \circ \gamma = \lambda \cdot \pi \circ \gamma = \lambda \cdot p \circ \gamma$$

is a rectifiable curve in V, its unique horizontal lift to $\mathbb{G}$ that starts at 0 is $\delta_\lambda \circ \gamma$. By the Lipschitz quotient property for p, this lift must be in $\widehat{X}$, so $\delta_\lambda(y) \in \widehat{X}$.

Thus, we find that $\widehat{X} = H$ is a homogeneous subgroup of $\mathbb{G}$ that is rectifiably connected. In particular, it is a closed Lie subgroup with Lie algebra $\mathfrak{h} \subset \mathfrak{g}$ for which $H = \exp_{\mathbb{G}}(\mathfrak{h})$. Moreover, the k–dimensional subspace $V = p(H) \subset V_1$ that we found earlier is precisely $\mathfrak{h} \cap V_1$. Let $\mathfrak{h}^* \subset \mathfrak{h}$ denote the stratified Lie algebra generated by V, and let $H^* = \exp_{\mathbb{G}}(\mathfrak{h}^*) \subset H$ be the corresponding Carnot subgroup of $\mathbb{G}$.

It remains only to show that $H^* = H$. To see this, let $y \in H$ be arbitrary, and let γ be a rectifiable curve in H from 0 to y. Then $p \circ \gamma$ is a rectifiable curve in V from 0 to $p(z)$, and its unique horizontal lift to $\mathbb{G}$, starting at 0, is γ. At the same time, $p \circ \gamma$ is a rectifiable curve in the horizontal layer of the Carnot subgroup H^*, so it has a horizontal lift to H^* through p that starts at 0. Hence, $y \in H^*$, and we conclude that $H^* = H$. □

The following is an immediate result of the theorem above:

Corollary 4.5 *Let (U, ϕ) be a k–dimensional chart in a complete, metrically doubling Lipschitz differentiability space (X, d, μ), and assume that for almost every $x \in U$, each $\widehat{X} \in \mathrm{Tan}(X, x)$ is rectifiably connected.*

Suppose that $F\colon X \to \mathbb{G}$ is a bi-Lipschitz embedding. Then for μ–almost every $x \in U$, $\mathrm{Tan}(F(X), F(x))$ consists of one element, which is canonically isometric to a Carnot

subgroup of $\mathbb{G}$ generated by a k–dimensional vector subspace of the horizontal layer of $\mathbb{G}$.

In particular, for μ–almost every $x \in U$, every element in $\mathrm{Tan}(X, x)$ is bi-Lipschitz equivalent to a sub-Riemannian Carnot group whose horizontal layer has dimension k.

4.1 Nonembedding consequences

One interesting consequence of Cheeger's initial study of differentiability in metric spaces was a certain generalized nonembedding result for Euclidean targets. The following statement was proven for PI spaces in [16, Theorems 14.1 and 14.2] and was generalized to Lipschitz differentiability spaces in [23, Corollary 8.3] as a corollary of Proposition 2.3 above.

Theorem 4.6 *Let (X, d, μ) be a complete, metrically doubling Lipschitz differentiability space with an k–dimensional chart (U, ϕ).*

Suppose there exists a set $A \subseteq U$ with $\mu(A) > 0$ such that for every $a \in A$, there exists $(Y, y) \in \mathrm{Tan}(X, a)$ that is not bi-Lipschitz equivalent to $\mathbb{R}^k$.

Then X does not admit a bi-Lipschitz embedding into any Euclidean space.

Cheeger and Kleiner [17, Theorem 1.6] used differentiability to prove a nonembedding result for PI spaces into certain infinite-dimensional Banach spaces.

Our work above allows us to obtain a result similar to Theorem 4.6 for Carnot group targets. Namely, a direct consequence of Corollary 4.5 is the following general nonembedding result:

Corollary 4.7 *Let (X, d, μ) be a complete, metrically doubling Lipschitz differentiability space with a k–dimensional chart (U, ϕ), and assume that for almost every $x \in U$, each $\widehat{X} \in \mathrm{Tan}(X, x)$ is rectifiably connected.*

Suppose there is a set $A \subseteq U$ with $\mu(A) > 0$ such that, for every $a \in A$, there exists $(Y, y) \in \mathrm{Tan}(X, a)$ that is not bi-Lipschitz equivalent to a Carnot group with a k–dimensional horizontal layer.

Then X does not admit a bi-Lipschitz embedding into any Carnot group.

A nonembedding result in the same spirit (but allowing more general targets) for a class of one-dimensional Lipschitz differentiability spaces already appears in [18, Corollary 9.3].

5 Rigidity for convex-cocompact actions of minimal Hausdorff dimension

This section is devoted to the proof of Theorem 1.4. It will be convenient to begin with the following lemma, which is essentially due to Bourdon, as part of the argument for Theorem 1.1.

Lemma 5.1 *Let $\Gamma \curvearrowright X$ be a convex-cocompact action of a discrete group Γ on a $\mathrm{CAT}(-1)$ space X. Suppose that the limit set $\Lambda(\Gamma)$ is bi-Lipschitz equivalent to the boundary of a noncompact rank-one symmetric space $S \neq \mathbb{H}^2_{\mathbb{R}}$. Then there is an isometric embedding $F\colon S \to X$ such that $F(\partial S) = \Lambda(\Gamma)$, and Γ stabilizes $F(S)$, acting on it with compact quotient.*

Proof Fix visual metrics on ∂X and ∂S once and for all. The bi-Lipschitz equivalence implies that $\dim_H(\Lambda(\Gamma)) = \dim_H(\partial S)$ and that the Möbius action $\Gamma \curvearrowright \Lambda(\Gamma)$ can be conjugated to a uniformly quasi-Möbius action $\Gamma \curvearrowright \partial S$.

If S is of quaternionic or octonionic type, then a theorem of Pansu [43, Corollary 11.2] ensures that the action $\Gamma \curvearrowright \partial S$ is by 1–quasiconformal mappings, hence mappings that extend to isometries of S by [43, Theorem 11.5], and hence Möbius mappings by [9].

If S is real or complex hyperbolic, then theorems of Sullivan and Tukia [47, Theorem G] and Chow [19, Theorem 2] show that there is a quasi-Möbius homeomorphism of ∂S that conjugates the quasi-Möbius action $\Gamma \curvearrowright \partial S$ to a Möbius action. To be more precise, here we use that every point in ∂S is a "radial limit point" for the quasi-Möbius action $\Gamma \curvearrowright \partial S$. This follows from the fact that this action is cocompact on triples, which in turn follows from the standard fact that the boundary action $\Gamma \curvearrowright \Lambda(\Gamma)$ is cocompact on triples. Also, in the complex hyperbolic case, we again turn to [43, Theorem 11.5; 9] to argue that conformal mappings are Möbius.

In any case, we can find a quasisymmetric homeomorphism $f\colon \partial S \to \Lambda(\Gamma)$ that is equivariant with respect to Möbius actions $\Gamma \curvearrowright \partial S$ and $\Gamma \curvearrowright \Lambda(\Gamma)$.

From here, Bourdon's work applies directly. In Section 2 of [10], he shows that such a homeomorphism f is in fact a Möbius homeomorphism. In addition, he shows in [10, Theorem 0.1] that any Möbius embedding $f\colon \partial S \to \partial X$ extends to an isometric embedding $F\colon S \to X$. That $\Gamma \curvearrowright X$ stabilizes $F(S)$ follows immediately from the fact that $F(S)$ is the union of geodesics in X whose endpoints are both in $\Lambda(\Gamma)$. That $F(S)/\Gamma$ is compact follows from the assumption that $\Gamma \curvearrowright X$ is convex-cocompact. □

Proof of Theorem 1.4 The inequality was already discussed in Remark 3.4, as was the equality case when $Q \leq 1$. Thus, we may assume that equality holds and $Q > 1$. It follows from Theorem 3.3 that $\Lambda(\Gamma)$, equipped with Q–dimensional Hausdorff measure, is a Lipschitz differentiability space. In fact, we know that $\Lambda(\Gamma)$ supports a Poincaré inequality, which means that it is a quasiconvex metric space (see Remark 4.2).

Let $\mathbb{G}$ be the Iwasawa group that locally models the geometry of $(\partial X, d_p)$. By Lemma 3.1, for any choice of $q \in X$ and $\omega \in \partial X$, the punctured boundary $\partial X \backslash \{\omega\}$ is identified with $\mathbb{G}$ in such a way that the parabolic visual metric $d_{\omega,q}$ is a homogeneous distance on $\mathbb{G}$.

Choose a point $z \in \Lambda(\Gamma) \subset \partial X$ for which every tangent of every element of the set $\mathrm{Tan}(\Lambda(\Gamma), d_p, z)$ is also in $\mathrm{Tan}(\Lambda(\Gamma), d_p, z)$. By [39, Theorem 1.1], almost every $z \in \Lambda(\Gamma)$, with respect to Hausdorff Q–measure, has this property. Take a tangent of $(\Lambda(\Gamma), d_p)$ at z. By [7, Proposition 3.1], there are points $q' \in X$ and $\omega' \in \Lambda(\Gamma)$ for which this tangent is isometric to the parabolic limit set $(\Lambda(\Gamma)\backslash\{\omega'\}, d_{\omega',q'})$. As $(\Lambda(\Gamma), d_p)$ is quasiconvex, its tangents are as well.

This parabolic limit set is a subset of the parabolic boundary $(\partial X\backslash\{\omega'\}, d_{\omega',q'})$, which is identified with the Iwasawa group $\mathbb{G}$. The metrics $d_{\omega',q'}$ and d_p are locally bi-Lipschitz equivalent on $\partial X\backslash\{\omega'\}$ and, hence, also on $\Lambda(\Gamma)\backslash\{\omega'\}$. In particular, this means that $(\Lambda(\Gamma)\backslash\{\omega'\}, d_{\omega',q'})$ is a Lipschitz differentiability space when equipped with its Hausdorff Q–measure. By Theorem 4.1, this parabolic limit set has a tangent that is isometric to a Carnot subgroup $\mathbb{N} \subset \mathbb{G}$.

By the way we chose $z \in \Lambda(\Gamma)$, we see that $(\mathbb{N}, d_{\omega',q'})$ is actually a tangent of the full limit set $(\Lambda(\Gamma), d_p)$. Once again, appealing to [7, Proposition 3.1], this means that there are points $q \in X$ and $\omega \in \Lambda(\Gamma)$ for which $(\mathbb{N}, d_{\omega',q'})$ is isometric to the parabolic limit set $(\Lambda(\Gamma)\backslash\{\omega\}, d_{\omega,q})$. Fix an identification of $\partial X\backslash\{\omega\}$ with $\mathbb{G}$ such that the origin of $\mathbb{G}$ is contained in $\Lambda(\Gamma)\backslash\{\omega\}$. In this way, $\mathbb{N} = \Lambda(\Gamma)\backslash\{\omega\}$ is a Carnot subgroup of $\mathbb{G} = \partial X\backslash\{\omega\}$, which is equipped with the metric $d_{\omega,q}$.

We claim that $\mathbb{N}$ is also an Iwasawa group. Suppose not, and let $g \in \Gamma$ be arbitrary. As remarked in Section 3.1, the induced boundary homeomorphism

$$g\colon \mathbb{G}\backslash\{g^{-1}(\omega)\} \to \mathbb{G}\backslash\{g(\omega)\}$$

is conformal, in the sense that it is smooth and its horizontal derivative is everywhere a similarity. In particular, the restriction

$$g\colon \mathbb{N}\backslash\{g^{-1}(\omega)\} \to \mathbb{N}\backslash\{g(\omega)\}$$

is also smooth with horizontal derivative everywhere a similarity, and so it is a conformal map on an open, connected subset of $\mathbb{N}$. By [22, Theorem 4.1], this means that $g|_{\mathbb{N}\setminus\{g^{-1}(\omega)\}}$ is the restriction of an affine map, and in particular that $g(\omega)=\omega$. Thus, ω is a global fixed point for the action of Γ on $\Lambda(\Gamma)$, a contradiction.

Let S be the noncompact rank-one symmetric space with ∂S locally modeled by the Iwasawa group $\mathbb{N}$. Note that S is determined uniquely by the topological and Hausdorff dimensions of $\mathbb{N}$, or equivalently of $\Lambda(\Gamma)$ [40, page 34]. Choosing any points $s\in S$ and $\zeta\in\partial S$, we know that $(\partial S\setminus\{\zeta\}, d_{\zeta,s})$ is bi-Lipschitz equivalent to $\mathbb{N}=(\Lambda(\Gamma)\setminus\{\omega\}, d_{\omega,q})$.

Using that $(\partial S\setminus\{\zeta\}, d_{\zeta,s})$ and $(\partial S\setminus\{\zeta\}, d_s)$ are Möbius equivalent, together with the fact that $(\Lambda(\Gamma)\setminus\{\omega\}, d_{\omega,q})$ and $(\Lambda(\Gamma)\setminus\{\omega\}, d_q)$ are Möbius equivalent, we find that $(\partial S\setminus\{\zeta\}, d_s)$ is quasi-Möbius equivalent to $(\Lambda(\Gamma)\setminus\{\omega\}, d_q)$ with linear distortion function $\eta(t)$. In particular, this implies that $(\partial S, d_s)$ and $(\Lambda(\Gamma), d_q)$ are bi-Lipschitz equivalent. As $\operatorname{Hdim}(\partial S)=\operatorname{Hdim}(\Lambda(\Gamma))=Q>1$, we know that $S\neq\mathbb{H}^2_{\mathbb{R}}$. Finally, applying Lemma 5.1 finishes the proof. □

References

[1] **D Bate**, *Structure of measures in Lipschitz differentiability spaces*, J. Amer. Math. Soc. 28 (2015) 421–482 MR

[2] **D Bate**, **S Li**, *Differentiability and Poincaré-type inequalities in metric measure spaces*, preprint (2015) arXiv

[3] **D Bate**, **S Li**, *Characterizations of rectifiable metric measure spaces*, Ann. Sci. Éc. Norm. Supér. 50 (2017) 1–37 MR

[4] **D Bate**, **G Speight**, *Differentiability, porosity and doubling in metric measure spaces*, Proc. Amer. Math. Soc. 141 (2013) 971–985 MR

[5] **S Bates**, **W B Johnson**, **J Lindenstrauss**, **D Preiss**, **G Schechtman**, *Affine approximation of Lipschitz functions and nonlinear quotients*, Geom. Funct. Anal. 9 (1999) 1092–1127 MR

[6] **M Bonk**, **B Kleiner**, *Rigidity for quasi-Möbius group actions*, J. Differential Geom. 61 (2002) 81–106 MR

[7] **M Bonk**, **B Kleiner**, *Rigidity for quasi-Fuchsian actions on negatively curved spaces*, Int. Math. Res. Not. 2004 (2004) 3309–3316 MR

[8] **M Bonk**, **B Kleiner**, *Conformal dimension and Gromov hyperbolic groups with* 2*–sphere boundary*, Geom. Topol. 9 (2005) 219–246 MR

[9] **M Bourdon**, *Structure conforme au bord et flot géodésique d'un* CAT(−1)*–espace*, Enseign. Math. 41 (1995) 63–102 MR

[10] **M Bourdon**, *Sur le birapport au bord des* CAT(−1)*–espaces*, Inst. Hautes Études Sci. Publ. Math. 83 (1996) 95–104 MR

[11] **M Bourdon**, **H Pajot**, *Poincaré inequalities and quasiconformal structure on the boundary of some hyperbolic buildings*, Proc. Amer. Math. Soc. 127 (1999) 2315–2324 MR

[12] **R Bowen**, *Hausdorff dimension of quasicircles*, Inst. Hautes Études Sci. Publ. Math. 50 (1979) 11–25 MR

[13] **M Bridgeman**, **R Canary**, **F Labourie**, **A Sambarino**, *The pressure metric for Anosov representations*, Geom. Funct. Anal. 25 (2015) 1089–1179 MR

[14] **S V Buyalo**, **A M Kuznetsov**, *Boundary at infinity of rank-one symmetric spaces*, Algebra i Analiz 21 (2009) 3–18 MR In Russian; translated in St. Petersburg Math. J. 21 (2010) 681–691

[15] **L Capogna**, **M Cowling**, *Conformality and Q–harmonicity in Carnot groups*, Duke Math. J. 135 (2006) 455–479 MR

[16] **J Cheeger**, *Differentiability of Lipschitz functions on metric measure spaces*, Geom. Funct. Anal. 9 (1999) 428–517 MR

[17] **J Cheeger**, **B Kleiner**, *Differentiability of Lipschitz maps from metric measure spaces to Banach spaces with the Radon–Nikodým property*, Geom. Funct. Anal. 19 (2009) 1017–1028 MR

[18] **J Cheeger**, **B Kleiner**, **A Schioppa**, *Infinitesimal structure of differentiability spaces, and metric differentiation*, Anal. Geom. Metr. Spaces 4 (2016) 104–159 MR

[19] **R Chow**, *Groups quasi-isometric to complex hyperbolic space*, Trans. Amer. Math. Soc. 348 (1996) 1757–1769 MR

[20] **C Connell**, *Minimal Lyapunov exponents, quasiconformal structures, and rigidity of non-positively curved manifolds*, Ergodic Theory Dynam. Systems 23 (2003) 429–446 MR

[21] **M Coornaert**, *Mesures de Patterson–Sullivan sur le bord d'un espace hyperbolique au sens de Gromov*, Pacific J. Math. 159 (1993) 241–270 MR

[22] **M G Cowling**, **A Ottazzi**, *Conformal maps of Carnot groups*, Ann. Acad. Sci. Fenn. Math. 40 (2015) 203–213 MR

[23] **G C David**, *Tangents and rectifiability of Ahlfors regular Lipschitz differentiability spaces*, Geom. Funct. Anal. 25 (2015) 553–579 MR

[24] **O Guichard**, *Groupes plongés quasi isométriquement dans un groupe de Lie*, Math. Ann. 330 (2004) 331–351 MR

[25] **U Hamenstädt**, *A geometric characterization of negatively curved locally symmetric spaces*, J. Differential Geom. 34 (1991) 193–221 MR

[26] **J Heinonen**, *Lectures on analysis on metric spaces*, Springer (2001) MR

[27] **J Heinonen**, **P Koskela**, *Quasiconformal maps in metric spaces with controlled geometry*, Acta Math. 181 (1998) 1–61 MR

[28] **J Heinonen**, **P Koskela**, **N Shanmugalingam**, **J T Tyson**, *Sobolev classes of Banach space-valued functions and quasiconformal mappings*, J. Anal. Math. 85 (2001) 87–139 MR

[29] **D Jerison**, *The Poincaré inequality for vector fields satisfying Hörmander's condition*, Duke Math. J. 53 (1986) 503–523 MR

[30] **W B Johnson**, **J Lindenstrauss**, **D Preiss**, **G Schechtman**, *Uniform quotient mappings of the plane*, Michigan Math. J. 47 (2000) 15–31 MR

[31] **I Kapovich**, **N Benakli**, *Boundaries of hyperbolic groups*, from "Combinatorial and geometric group theory" (S Cleary, R Gilman, A G Myasnikov, V Shpilrain, editors), Contemp. Math. 296, Amer. Math. Soc., Providence, RI (2002) 39–93 MR

[32] **S Keith**, *A differentiable structure for metric measure spaces*, Adv. Math. 183 (2004) 271–315 MR

[33] **S Keith**, **T Laakso**, *Conformal Assouad dimension and modulus*, Geom. Funct. Anal. 14 (2004) 1278–1321 MR

[34] **S Keith**, **X Zhong**, *The Poincaré inequality is an open ended condition*, Ann. of Math. 167 (2008) 575–599 MR

[35] **K Kinneberg**, *Rigidity for quasi-Möbius actions on fractal metric spaces*, J. Differential Geom. 100 (2015) 349–388 MR

[36] **B Kleiner**, **J M Mackay**, *Differentiable structures on metric measure spaces: a primer*, Ann. Sc. Norm. Super. Pisa Cl. Sci. 16 (2016) 41–64 MR

[37] **A Korányi**, **H M Reimann**, *Quasiconformal mappings on the Heisenberg group*, Invent. Math. 80 (1985) 309–338 MR

[38] **R Korte**, **N Marola**, **N Shanmugalingam**, *Quasiconformality, homeomorphisms between metric measure spaces preserving quasiminimizers, and uniform density property*, Ark. Mat. 50 (2012) 111–134 MR

[39] **E Le Donne**, *Metric spaces with unique tangents*, Ann. Acad. Sci. Fenn. Math. 36 (2011) 683–694 MR

[40] **J M Mackay**, **J T Tyson**, *Conformal dimension: theory and application*, University Lecture Series 54, Amer. Math. Soc., Providence, RI (2010) MR

[41] **R Montgomery**, *A tour of subriemannian geometries, their geodesics and applications*, Mathematical Surveys and Monographs 91, Amer. Math. Soc., Providence, RI (2002) MR

[42] **P Pansu**, *Dimension conforme et sphère à l'infini des variétés à courbure négative*, Ann. Acad. Sci. Fenn. Ser. A I Math. 14 (1989) 177–212 MR

[43] **P Pansu**, *Métriques de Carnot–Carathéodory et quasiisométries des espaces symétriques de rang un*, Ann. of Math. 129 (1989) 1–60 MR

[44] **A Schioppa**, *Derivations and Alberti representations*, preprint (2013) arXiv

[45] **A Schioppa**, *Derivations and Alberti representations*, Adv. Math. 293 (2016) 436–528 MR

[46] **K Tan**, **X Yang**, *Subriemannian geodesics of Carnot groups of step* 3, ESAIM Control Optim. Calc. Var. 19 (2013) 274–287 MR

[47] **P Tukia**, *On quasiconformal groups*, J. Analyse Math. 46 (1986) 318–346 MR

[48] **J Väisälä**, *Quasi-Möbius maps*, J. Analyse Math. 44 (1984/85) 218–234 MR

GCD: *Department of Mathematical Sciences, Ball State University*
Muncie, IN, United States

GCD: *Courant Institute of Mathematical Sciences, New York University*
New York, NY

KK: *National Security Agency*
Fort Meade, MD, United States

KK: *Department of Mathematics, Rice University*
Houston, TX, United States

gcdavid@bsu.edu, kekinneberg@gmail.com

https://sites.google.com/view/gcdavid,
https://sites.google.com/site/kekinneberg

Proposed: Bruce Kleiner
Seconded: Benson Farb, John Lott

Received: 15 September 2016
Revised: 26 January 2018