

BiLipschitz homogeneous hyperbolic nets

CHRISTOPHER J. BISHOP

Abstract. We answer a question of Itai Benjamini by showing there is a $K < \infty$ so that for any $\epsilon > 0$, there exist ϵ -dense discrete sets in the hyperbolic disk that are homogeneous with respect to K -biLipschitz maps of the disk to itself. However, this is not true for K close to 1; in that case, every K -biLipschitz homogeneous discrete set must omit a disk of hyperbolic radius $\epsilon(K) > 0$. For $K = 1$, this is a consequence of the Margulis lemma for discrete groups of hyperbolic isometries.

Kaksisuuntaisten Lipschitzin kuvausten suhteen tasalaatuista hyperbolisista verkoista

Tiivistelmä. Tässä työssä vastataan Itai Benjaminin esittämään kysymykseen osoittamalla, että on olemassa sellainen $K < \infty$, että jokaista lukua $\epsilon > 0$ kohti on olemassa hyperbolisen kiekon ϵ -tiheä diskreetti joukko, joka on tasalaatuinen kiekon itselleen kuvaavien, kaksisuuntaisen Lipschitzin ehdon vakiolla K toteuttavien kuvausten suhteen. Tämä ei kuitenkaan päde, jos luku K on lähellä ykköstä; siinä tapauksessa jokainen diskreetti joukko, joka on tasalaatuinen em. kuvausten suhteen, välttämättä väistää jonkin kiekon, jonka hyperbolinen säde on $\epsilon(K) > 0$. Kun $K = 1$, tämä seuraa hyperbolisten isometrioiden diskreettejä ryhmiä koskevasta Marguliksen lemmasta.

1. Introduction

Let $\mathbb{D} = \{z: |z| < 1\}$ denote the unit disk in the complex plane $\mathbb{C}$ and let ρ denote the hyperbolic metric on $\mathbb{D}$ (defined in Section 2). A set $X \subset \mathbb{D}$ is called discrete if it has no accumulation points in $\mathbb{D}$, and for $\epsilon > 0$ it is called ϵ -dense if every $z \in \mathbb{D}$ is within hyperbolic distance ϵ of some point $x \in X$. A set X is called homogeneous with respect to a set $\mathcal{F}$ of homeomorphisms if for any $x, y \in X$ there is a $f \in \mathcal{F}$ so that $f(X) = X$ and $f(x) = y$. In other words, $\mathcal{F}$ acts transitively on X (we do not assume $\mathcal{F}$ is a group; see Remark 1 below). We say that $X \subset \mathbb{D}$ is a (K, ϵ) -net if it is a discrete ϵ -net that is homogeneous with respect to the set of hyperbolic K -biLipschitz maps from $\mathbb{D}$ onto itself (we could also consider biLipschitz self-maps of X ; see Remark 2 below). Define

$$\epsilon(K) = \inf\{\epsilon: (K, \epsilon)\text{-nets exist}\}.$$

This is finite for all $K \geq 1$ since it is clearly a decreasing function of K (as K increases, the infimum is over larger sets), and the orbit of any co-compact Fuchsian group G is a $(1, \epsilon)$ -net for some $\epsilon < \infty$; we can take ϵ to be the diameter of the compact quotient surface $R = \mathbb{D}/G$. An explicit bound is given by the genus two Bolza surface, whose hyperbolic diameter is $\arctan(3 + 2\sqrt{2}) \approx 2.45$; see [10]. Thus

$$K_c = \inf\{K: \epsilon(K) = 0\} = \sup\{K: \epsilon(K) > 0\}$$

is well defined and $1 \leq K_c \leq \infty$. We shall prove both inequalities are strict.

<https://doi.org/10.54330/afm.152404>

2020 Mathematics Subject Classification: Primary 30F45; Secondary 30F10, 51M09.

Key words: BiLipschitz maps, hyperbolic geometry, Margulis constant, homogeneous set, quadrilateral mesh.

The author is partially supported by NSF Grant DMS 2303987.

© 2024 The Finnish Mathematical Society

Theorem 1.1. $1 < K_c < \infty$.

The upper bound follows from an explicit construction: for any $\epsilon > 0$ we build an ϵ -net that is homogeneous for K -biLipschitz maps with $K < \infty$ independent of ϵ . The lower bound is given by an indirect argument. Assuming $K_c = 1$ we show there exists an ϵ -dense set in $\mathbb{D}$ that “looks like” a copy of $\mathbb{Z} \times \mathbb{Z}$, and we will derive a contradiction with the exponential growth of the hyperbolic area.

It is well known that $\epsilon(1) > 0$. If X is homogeneous with respect to hyperbolic isometries, then these maps generate a subgroup H of the group G of all hyperbolic isometries mapping X to itself. Since X is a discrete set, G is a discrete group, i.e., a Fuchsian group, and $R = \mathbb{D}/G$ is a (possibly branched) Riemann surface and the set X projects to a single point $x \in R$. A famous result of Kařdan and Margulis [7], says that there is a positive constant $\epsilon_1 > 0$ (the Margulis constant) so that the injectivity radius is at least ϵ_1 at some point of R , and hence R contains a disk of radius at least $\epsilon_1/2$ that does not intersect X . Thus $\epsilon(1) \geq \epsilon_1/2$. Alternate proofs of the Margulis lemma for Fuchsian groups are given in [8, 11, 13]; the latter gives the sharp value. The question of whether $K_c > 1$ was raised by Itai Benjamini as a result of considering whether the Margulis lemma really requires the machinery of hyperbolic isometries, group actions and fundamental domains, or might it have an analog for sets of biLipschitz mappings.

Remark 1. We claim that $K_c = \infty$, if we require X to be homogeneous with respect to some group H of K -biLipschitz maps on $\mathbb{D}$. Such a group would consist of K^2 -quasiconformal maps, and a result of Tukia [12] says that such a group is of the form $H = hGh^{-1}$ for some quasiconformal map $h: \mathbb{D} \rightarrow \mathbb{D}$ and some Möbius group G acting on $\mathbb{D}$. By Mori’s theorem [9] (or [1, Chapter 3]), the image of a hyperbolic ϵ -disk under h or h^{-1} contains a hyperbolic disk of radius $\geq \epsilon^{K^2}/16$. Since $h(X)$ is invariant under G , the previous paragraph shows it omits some disk of hyperbolic radius ϵ_1 , and hence X omits some disk of radius ϵ_K depending only on K .

Remark 2. We have assumed that X is homogeneous under biLipschitz self-maps of the disk, but we could replace this by self-maps of X . Our proof of the upper bound produces biLipschitz maps of the whole disk, and the proof of the lower bound only uses that we have self-maps of X . Thus the inequality $1 < K_c < \infty$ holds in either case, although it is not clear whether the exact value of K_c is the same in both situations (this depends on whether a K -biLipschitz self-map of an ϵ -net can be extended to a K -biLipschitz self-map of the disk, or whether a larger constant is sometimes needed).

I thank the anonymous referee for a careful reading of the manuscript and several suggestions that improved it.

2. The upper bound: $K_c < \infty$

The pseudo-hyperbolic metric on $\mathbb{D}$ is given by

$$\tilde{\rho}(z, w) = \left| \frac{z - w}{1 - \bar{w}z} \right|$$

and the hyperbolic metric by

$$\rho(z, w) = \frac{1}{2} \log \frac{1 + \tilde{\rho}(z, w)}{1 - \tilde{\rho}(z, w)}.$$

The hyperbolic metric can also be defined as $\rho(z, w) = \inf \int_{\gamma} ds / (1 - |x|^2)$, where the infimum is over all rectifiable paths in $\mathbb{D}$ connecting z and w . This implies $\rho(z, w) > |z - w|$ whenever $z \neq w$. The (orientation preserving) isometries of the hyperbolic metric are the linear fractional transformations of the disk to itself. The geodesics for the hyperbolic metric are diameters of the circle and their images under isometries, i.e., circular arcs perpendicular to the boundary. A ball of hyperbolic radius r has hyperbolic area that grows exponentially in r . See [2] or [5] for these basic facts about the hyperbolic metric. A hyperbolic K -biLipschitz map $f: X \rightarrow Y$ between subsets of $\mathbb{D}$ is one that satisfies

$$1/K \leq \frac{\rho(f(z), f(w))}{\rho(z, w)} \leq K \quad \text{for all } z, w \in X.$$

In this section, we prove the upper bound $K_c < \infty$ in Theorem 1.1 by building explicit (K, ϵ) -nets with K fixed and ϵ tending to zero. All our examples correspond to infinite quadrilateral meshes that refine a fixed tessellation of $\mathbb{D}$ by right pentagons. These meshes were constructed for different purposes in [4] (in that paper, they are part of the proof that any simple planar n -gon can be quad-meshed in time $O(n)$ using elements with all new angles between 60° and 120°).

Lemma 2.1. (Quadmeshes exist) *For all sufficiently small $\epsilon > 0$, the hyperbolic unit disk has a mesh by quadrilaterals each containing, and contained in, disks of size comparable to ϵ .*

Proof. We start with the standard tessellation of $\mathbb{D}$ by hyperbolic right pentagons. See the left side of Figure 1. Connect the center of each pentagon to the midpoint of each of its five boundary arcs. This divides the pentagon into five fundamental quadrilaterals. Each quadrilateral has three right angles and one angle of $2\pi/5$, the latter at the center of the pentagon.




Figure 1. Hyperbolic right pentagons tessellate the disk. Each pentagon is divided into five quadrilaterals which are then each divided into a $N \times N$ quadrilateral mesh (here $N = 3$). The elements all have hyperbolic diameter and side lengths $\simeq 1/N$.

The two edges of such a quadrilateral Q that are adjacent to the center of the pentagon have the same hyperbolic length as each other, as do the two sides of Q that are opposite these sides. Choose a positive integer N and divide each quadrilateral into a $N \times N$ quadrilateral mesh using geodesic arcs as shown in the center of Figure 1. Each boundary arc of the fundamental quadrilateral is divided into N sub-arcs of equal length. This implies the mesh in each fundamental quadrilateral matches the mesh in all its neighbors and defines a quadrilateral mesh of the whole disk. See the right side of Figure 1. We will call this mesh M ; it is an infinite graph embedded in $\mathbb{D}$ in which every vertex has degree four or five (the latter occurs only at the centers

of the hyperbolic pentagons). The set X of vertices of this mesh is our ϵ -net with $\epsilon \simeq 1/N$. $\square$

We observe the following for later use.

Lemma 2.2. *Suppose M is one of the meshes constructed above. The hyperbolic distance between two vertices z, w of M is comparable to their graph distance, $d_M(w, z)$ divided by N , with a constant that is independent of M .*

Proof. To see this, note that each edge of the mesh has hyperbolic length $O(1/N)$, so $\rho(z, w) = O(d_M(z, w)/N)$. On the other hand, a geodesic segment $\gamma \subset \mathbb{D}$ connecting distinct points $z, w \in X$ can hit at most $O(N\rho(z, w))$ faces of the mesh: each such face has hyperbolic area $\simeq 1/N^2$, and is contained in a $O(1/N)$ neighborhood of γ , so the union of faces hitting γ has area $O(\rho(z, w)/N)$. The edges of these faces contain a path of mesh edges connecting z and w , so $d_M(z, w) = O(N\rho(z, w))$. $\square$

We let $Y \subset X$ denote the vertices of the pentagonal tessellation in the left side of Figure 1; these are the points where the geodesics defining the edges of the tessellation cross each other. We call these geodesic edges the “bounding geodesics” and call their crossing points Y the “corner points”. The set of corner points is clearly homogeneous under isometries of the hyperbolic disk. Thus to map a point $x_1 \in X$ to another point $x_2 \in X$, it suffices to map x_1 to some $y_1 \in Y$ and map some $y_2 \in Y$ to x_2 , and then isometrically map y_1 to y_2 . Thus it is enough to show that each x inside a fundamental quadrilateral Q can be mapped to a corner point $y \in Y$ by a K -biLipschitz map of X to itself, with K independent of N .

We do this in two steps. Given $x \in X$, let Q be the fundamental quadrilateral containing x and let $y = Q \cap Y$ be the corresponding corner point. First we will define a “discrete rotation” of X around y that maps x to a point $z \in \partial Q \cap X$ that lies on bounding geodesic γ passing through y . The second step is to define a “discrete translation” of X along γ that maps z to y . We will show both steps can be accomplished by K -biLipschitz maps, with K independent of N .

If $x \in Y$, there is nothing to do, so we assume $x \notin Y$ and choose $y = Q \cap Y$ where Q is a fundamental quadrilateral containing x . If x is on a bounding geodesic passing through y , we can continue to the second step of the construction, so for the moment, we assume this is not the case.

Lemma 2.3. (Discrete rotations) *Assume notation is as above. There is a $K < \infty$ so that the following holds. Suppose Q is a fundamental quadrilateral and $y \in Y$ is a corner point of Q . If $x \in X \cap Q$, then there is a K -bilipschitz map of the hyperbolic disk to itself that maps X to itself and maps x to a point $z \in X$ on a bounding geodesic γ that passes through y . The map is the identity off the four tessellation pentagons touching y .*

Proof. The corner point y is on the boundary of four hyperbolic pentagons. Let P be the union of these four pentagons. We define a series of closed cycles $\{\Gamma_k\}_1^{2N}$ in the mesh $M \cap P$. See Figure 2 for an example where $N = 5$. The first curve, Γ_1 , consists of the eight points of X that are adjacent to y in the mesh M . In general, if we have already defined $\Gamma_1, \dots, \Gamma_k$, then Γ_{k+1} is the cycle consisting of points of $X \cap P$ that are adjacent to Γ_k but not in Γ_{k-1} . Note that Γ_{2N} lies on the boundary of P . Also, for $k = 1, \dots, N$, observe that Γ_{k+1} has eight more points than Γ_k . For $k \geq N$, Γ_{k+1} has sixteen more points than Γ_k (note that Γ_N is the cycle passing through the centers of the pentagons).

By assumption, $x \neq y$, but x and y are in the same fundamental quadrilateral Q , so x lies on some Γ_k with $1 \leq k \leq N$. Moreover, there is a point $w \in X \cap \Gamma_k$ that is on a bounding geodesic and is at most $j \leq k$ steps away from x on the cycle Γ_k . Thus we can map x to w by “rotating” Γ_k by j steps (every point of Γ_k is moved j positions in the same direction).




Figure 2. The left shows the union of four pentagons (= twelve fundamental quadrilaterals) that touch the corner point y at the center. The right picture shows the concentric cycles Γ_k surrounding y . If x is on the k th ring with $1 \leq k \leq N$, then it is at most k points away from a bounding geodesic passing through y (illustrated here by the vertical and horizontal lines through y). Here $N = 5$.

We extend this rotation to the rest of X as follows. For $1 \leq m < j$ we rotate Γ_{k+m} by $j - |m|$ positions. Similarly for Γ_{k-m} . On the rest of X we take the identity. Recall that the hyperbolic distance between points of X is comparable (with absolute constants) to the mesh distance in M divided by N . The map above clearly only multiplies mesh distances by at most a bounded factor, independent of N . To see this, note that if $m \neq 0$ and two points are on Γ_k and Γ_{k+m} respectively, then the mesh distance between them is at least m and it can increase by at most $O(|m|)$ (since the size of the shifts varying by at most m , and partly due to the lengths of the two cycles differing by at most $16m$). If two points are on the same cycle Γ_k , then the shifts at worst multiply the mesh distance by two. Since $d_M(z, w) \simeq N\rho(z, w)$, our maps also multiply hyperbolic distances by a bounded factor, i.e. they are hyperbolically Lipschitz with a uniform constant. Moreover, the inverse map has the same form, so the inverse is also Lipschitz with a uniform bound. Thus our discrete rotation map is uniformly biLipschitz as a map $X \rightarrow X$.

We can extend the map $X \rightarrow X$ defined above to be a biLipschitz self-map of the whole disk. We define the extension as the identity outside P (the union of pentagons touching y), and within P we define it as follows. For each annular region A_k between the cycles Γ_{k-1} and Γ_k we take a biLipschitz map of A_k to a round annulus with points of X mapping to evenly spaced points on each boundary circle (this is easy). The discrete rotation maps on the cycles become Euclidean rotations on the boundary circles, and the angle of rotations differ by at most a bounded multiple of the width of the annulus. These boundary rotations can be interpolated by a biLipschitz map that just rotates each concentric circle between the boundaries, and this map is then transported back to A_k . $\square$

This completes the first step of our construction: every $x \in X$ can be mapped to a point of X lying on a bounding geodesic of the pentagonal tessellation, by a

uniformly biLipschitz map of $\mathbb{D}$ to itself. Next we need to show any such mesh point on a bounding geodesic can be mapped to a corner point $y \in Y$. This is easier than the previous step.




Figure 3. S (for strip) is the union of fundamental quadrilaterals touching a single bounding geodesic (the thickened central horizontal line passing through y , the white dot). The mesh in this region is isomorphic to the square mesh on $\mathbb{Z} \times [-N, N]$ and we can define a biLipschitz map that translates the central geodesic by j mesh elements (with $|j| \leq N$) and is the identity outside S . Here $N = 5$.

Lemma 2.4. (Discrete translations) *There is a $K < \infty$ so that the following holds. Suppose, as above, that Q is a fundamental quadrilateral, y is a corner point of Q , γ is bounding geodesic for Q , and $z \in \gamma \cap Q$. Then there is a K -bilipschitz map of the hyperbolic disk to itself that maps X to itself and maps z to y . The map is the identity off the union of fundamental quadrilaterals that touch γ .*

Proof. Let S denote the union of all fundamental quadrilaterals Q that touch γ . In Figure 3 a single fundamental quadrilateral Q is shaded. This quadrilateral, and its three rotations around y by $\pi/2$, π and $3\pi/2$, form a larger quadrilateral Q' , and S is union of translates of Q' under powers (positive and negative) of a single hyperbolic translation along γ . See Figure 3.

The restriction of the mesh M to S is isomorphic to the graph $\mathbb{Z} \times [-N, N]$. Fixing j with $1 \leq j < N$, we can define a discrete translation by j by sending $(n, m) \rightarrow (n + j - |m|, m)$ for $|m| < j$ and taking the identity map on the rest of X . This is clearly uniformly biLipschitz in both the graph and hyperbolic distances on X and can move a point on a bounding geodesic up to N positions. This is enough to move any point onto a corner, as desired. The extension to a biLipschitz self-map of the disk is similar to, but even simpler than, the previous case. We take the identity map outside S . We map the strip S to the Euclidean strip $\mathbb{R} \times [-N, N]$, extend the translations on the top and bottom edges to the interior via a shear map preserving each horizontal line, and then map this back to S . $\square$

This completes the proof of upper bound in Theorem 1.1.

3. The lower bound: $K_c > 1$

In this section, we prove $K_c > 1$ by contradiction. We will assume that $K_c = 1$ and construct an ϵ -dense mesh that only has $O(r/\epsilon^2)$ points within hyperbolic distance r of the origin. However, this contradicts the well known fact that a hyperbolic ball of radius r has area that grows exponentially with r . In this section, we need only assume that X is homogeneous with respect to K -biLipschitz homeomorphisms of X to itself.

Lemma 3.1. *If X is a (K, ϵ) -net, then it is also a (K, δ) net, where $\delta \leq \epsilon$ is the supremum of numbers $r > 0$ so that $\mathbb{D} \setminus X$ contains a disk of radius r .*

Proof. Clearly $\delta > 0$, since X discrete and hence its complement is open and non-empty. Also $\delta \leq \epsilon$, since every point is within ϵ of some point of X . The set X is a (K, δ) -net, since every point of $\mathbb{D}$ is within hyperbolic distance δ of X by definition. $\square$

Henceforth, given an (K, ϵ) -net, we assume we have set $\epsilon = \delta$ as above. Thus for any $\lambda \in (0, 1)$, there is a point $x_\lambda \in X$ and a hyperbolic disk D_λ of radius $\lambda\epsilon$ so that $D_\lambda \cap X = \emptyset$, and whose center is within distance ϵ of x_λ .

Lemma 3.2. *Suppose that X is homogeneous with respect to K -biLipschitz homeomorphisms of X to itself. If $K > 1$ is close enough to 1 and $\epsilon > 0$ is close enough to 0, then every closed disk of radius ϵ around any point $w \in X$ contains the center of a disk of radius $\epsilon/4$ that is not hit by X .*

Proof. First, note that the regular 12-gon inscribed in a Euclidean disk of radius 1 has side lengths $2\sin(\pi/12) \approx .5176$, so that in hyperbolic space 12 equally spaced points on a circle of hyperbolic radius ϵ will be more than $2\sin(\pi/12)\epsilon > \epsilon/2$ apart. This is because on small neighborhoods of the origin, hyperbolic distances approximate Euclidean distances. Now place twelve disjoint disks of hyperbolic radius $\epsilon/4$ on the hyperbolic circle of radius ϵ around point $w \in \mathbb{D}$. By our assumption that the conclusion of the lemma fails, each of these twelve disks contains a point $\{x_k\}_{k=1}^{12} \subset X$. See Figure 4.




Figure 4. In hyperbolic space we can place twelve disks of radius $\epsilon/4$ on a circle of radius ϵ around $w \in X$ (black dot in center). Under our assumptions, each smaller disk of radius $\epsilon/4$ contains a point of X (white dots).

Now map w to x_λ by a K -biLipschitz map of X . The distances between w and the twelve images of $\{x_k\}_1^{12}$, and among these twelve points change by at most a factor of $K \approx 1$. This determines the positions of the twelve image points around x_λ up to a rotation around x_λ , and an error that is $o(\epsilon)$ as K tends to 1.

Moreover, as ϵ tends to zero, the length of the gaps on circle $\{z: |z - w| = \epsilon\}$ between these evenly spaced disks tends to $\epsilon(2\pi - 24\sin(\pi/12))/12 \approx (.006)\epsilon$. Thus every point of the circle of radius ϵ around w is within $3\epsilon/4$ of some point of $\{x_k\}_1^{12}$. So if K is close enough to 1, then every point on the circle of radius ϵ around x_λ is within $7\epsilon/8$ of an image of one of the twelve points. If we choose $\lambda > 15/16$, then the center of D_λ is less than $\epsilon/16$ from this circle, and hence it is within $\lambda\epsilon$ of the center of D_λ . Thus D_λ contains a point of X , a contradiction. $\square$

Lemma 3.3. $K_c > 1$.

Proof. Suppose this fails. Then we can find sequences of sets $\{X_n\}$ and numbers $K_n \searrow 1$ and $\epsilon_n \searrow 0$ so that X_n is a (K_n, ϵ_n) -net. We will show that for n large enough, the sets X_n must have a local Euclidean structure that is incompatible with their global hyperbolic structure.

Restrict each X_n to the Euclidean disk $D(0, \sqrt{\epsilon_n}) \subset \mathbb{D}$ and expand it by the Euclidean dilation $z \rightarrow z/\epsilon_n$. This rescaling gives a sequence of finite sets $Z_n \subset \mathbb{C}$, from which we can extract a sequence that converges (in the Hausdorff metric on compact subsets of the plane) to closed set Z in $\mathbb{R}^2$ that satisfies:

- (1) Z is a (Euclidean) 1-net,
- (2) Z is homogeneous with respect to Euclidean isometries,
- (3) Any 2-ball centered in Z contains a $\frac{1}{4}$ -ball disjoint from Z .

The set of isometries that map Z into itself is a closed subgroup G of the Euclidean isometry group. Since G acts transitively on the 1-net Z , it must be infinite. Thus G is a closed, infinite Lie subgroup of the isometry group of the plane and hence must be either a discrete group (in which case, Z is a Euclidean lattice) or $G = \mathbb{R} \times \mathbb{Z}$ (and Z is a union of evenly spaced parallel lines).

In either case, Z contains a lattice Z' whose fundamental parallelogram is close to a square. In the case $Z = \mathbb{R} \times \mathbb{Z}$, we can take an actual square sub-lattice, and otherwise we can choose elements of Z that are within distance 1 of the points 10 and $10i$; these give a fundamental parallelogram of uniformly bounded eccentricity (all angles bounded uniformly away from 0 and π).

This means if ϵ_n and K_n are close enough to 0 and 1 respectively, then setting $\delta = 10\epsilon$ and taking any point $x \in X_n$ we can find eight other points in X_n that approximate a Euclidean 3×3 lattice with side lengths comparable to δ centered at x . Applying the same argument to each of the eight boundary points of this grid, we can expand it to a 5×5 grid. Continuing, we can build a $(2m+1) \times (2m+1)$ grid centered at x that is a union of approximate δ sized quadrilaterals that approximate squares uniformly. More precisely, we obtain a subset of X that is a 2δ -net in $\mathbb{D}$ whose points are $\delta/2$ separated, and has the structure of a Euclidean square mesh. But then $O(m^2)$ disks of radius 2δ centered on this grid cover a ball of radius $\simeq \delta n$ around x . In other words, $O(m^2)$ disks, each of hyperbolic area $O(\delta^2)$, cover a ball of hyperbolic radius $\simeq m\delta$, and this ball must have hyperbolic area at least $\exp(c\delta m)$ for some $c > 0$. This is impossible for large m , and the contradiction implies $K_c > 1$. $\square$

4. Questions and remarks

What is the precise value of K_c ?

Is the function $\epsilon(K)$ strictly decreasing on $[1, K_c]$? Is $\epsilon(K_c) = 0$? Does $\epsilon(K)$ tend to the Margulis constant as $K \searrow 1$? Is $\epsilon(K)$ continuous? It seems possible that the nets that minimize ϵ for a given K could have some special combinatorial structure, and that when this is changed, the optimal ϵ is different. Thus it seems possible that jumps in $\epsilon(K)$ could occur.

What can happen if X is a K -biLipschitz ϵ -net, but we don't require X be discrete? Then we could have $X = \mathbb{D}$; what else is possible? In general, a K -biLipschitz homogeneous compact set in $\mathbb{R}^2$ can be a Cantor set, even with K close to 1 (think of a thin Cantor set constructed using very thick annuli; the outer and inner boundary boundaries can be rotated all the way around by a biLipschitz map with small constant). What if X has non-trivial connected components? Hoehn

and Oversteegen [6] proved any compact planar set that is homogeneous under self-homeomorphisms is necessarily either a finite set, a Cantor set, a Jordan curve, a pseudo-arc, a circle of pseudo-arcs or the product of one of the first two with one of the latter three. It is still unknown (at least to the author) whether a biLipschitz homogeneous continuum in $\mathbb{R}^2$ must be a Jordan curve. However, it is known that a biLipschitz homogeneous Jordan curve in the plane must be a quasicircle [3].

The referee of this paper asked if analogous results hold in higher dimensions. It seems quite possible that both the explicit construction of (K, ϵ) -nets used to prove $K_c < \infty$, and the rescaling argument that proves $K_c > 1$ should work in higher dimensional hyperbolic spaces, but the details are not obvious. For example, what is the precise analog of the pentagonal tessellation of the disk that we used, and what are the analogs of the discrete rotations and translation maps? Assuming an extension of our result to higher dimensions is possible, what is the behavior of K_c as the dimension increases to infinity? The proof given in this paper that $K_c > 1$ depended on the observation that hyperbolic space looks Euclidean at small scales, but not at large scales. Does the result in this paper extend to other (non-hyperbolic) spaces that are “non-Euclidean” at large scales?

References

- [1] AHLFORS, L. V.: Lectures on quasiconformal mappings. - Univ. Lecture Ser. 38, Amer. Math. Soc., Providence, RI, second edition, 2006.
- [2] BEARDON, A. F.: The geometry of discrete groups. - Grad. Texts in Math. 91, Springer-Verlag, New York, 1983.
- [3] BISHOP, C. J.: Bi-Lipschitz homogeneous curves in $\mathbb{R}^2$ are quasicircles. - Trans. Amer. Math. Soc. 353:7, 2001, 2655–2663.
- [4] BISHOP, C. J.: Optimal angle bounds for quadrilateral meshes. - Discrete Comput. Geom. 44:2, 2010, 308–329.
- [5] GARNETT, J. B., and D. E. MARSHALL: Harmonic measure. - New Math. Monogr. 2, Cambridge Univ. Press, Cambridge, 2008.
- [6] HOEHN, L. C., and L. G. OVERSTEEGEN: A complete classification of homogeneous plane continua. - Acta Math. 216:2, 2016, 177–216.
- [7] KAŽDAN, D. A., and G. A. MARGULIS: A proof of Selberg’s hypothesis. - Mat. Sb. (N.S.) 75 (117), 1968, 163–168.
- [8] MARDEN, A.: Universal properties of Fuchsian groups in the Poincaré metric: - In: Discontinuous groups and Riemann surfaces (Proc. Conf., Univ. Maryland, College Park, Md., 1973), Ann. of Math. Stud. 79, 1974, 315–339.
- [9] MORI, A.: On an absolute constant in the theory of quasi-conformal mappings. - J. Math. Soc. Japan 8, 1956, 156–166.
- [10] STEPANYANTS, H., A. F. BEARDON, J. PATON, and D. KRIOUKOV: Diameter of compact Riemann surfaces. - arXiv:2301.10844v3 [math.GT], 2024.
- [11] STURM, J., and M. SHINNAR: The maximal inscribed ball of a Fuchsian group. - In: Discontinuous groups and Riemann surfaces (Proc. Conf., Univ. Maryland, College Park, Md., 1973), Ann. of Math. Stud. 79, 1974, 439–443.
- [12] TUKIA, P.: On two-dimensional quasiconformal groups. - Ann. Acad. Sci. Fenn. Ser. A I Math. 5:1, 1980, 73–78.
- [13] YAMADA, A.: On Marden’s universal constant of Fuchsian groups. - Kodai Math. J. 4:2, 1981, 266–277.

Received 27 March 2024 • Revision received 6 October 2024 • Accepted 7 October 2024
Published online 26 November 2024

Christopher J. Bishop
Stony Brook University
Mathematics Department
Stony Brook, NY 11794-3651, USA
bishop@math.stonybrook.edu