Annales Academiæ Scientiarum Fennicæ
Mathematica
Volumen 30, 2005, 313–336

CHARACTERIZATIONS OF SNOWFLAKE METRIC SPACES

Jeremy T. Tyson and Jang-Mei Wu

University of Illinois, Department of Mathematics
1409 West Green Street, Urbana, IL 61822, U.S.A.: tyson@math.uiuc.edu, wu@math.uiuc.edu

Abstract. A metric space (X,d) is said to be an L^p-metric space, $p\in[1,\infty)$, if $d(x,y)^p\le d(x,z)^p+d(z,y)^p$ for all $x,y,z\in X$, and is said to be a snowflake if it is bi-Lipschitz equivalent with an L^p-metric space for some $p>1$. Suppose that (X,d) is compact and doubling. Then (X,d) is a snowflake if and only if X admits a bi-Lipschitz embedding in a uniformly convex Banach space and no weak tangent to (X,d) contains a rectifiable curve. We give several equivalent conditions for the snowflake property, and examples distinguishing these conditions under weaker hypotheses. As a corollary we prove that the polygasket $\mathrm{PG}(N)\subset\mathbf{R}^2$ is a snowflake for $N=5$ or $N\ge 7$.

1. Introduction

Fix $p>1$. A metric space (X,d) is called a p-*snowflake* if there is a metric d', bi-Lipschitz equivalent with d, so that the L^p triangle inequality

$$d'(x,y)\le\begin{cases}\{d'(x,z)^p+d'(z,y)^p\}^{1/p}, & p<\infty,\\ \max\{d'(x,z),d'(z,y)\}, & p=\infty,\end{cases}\tag{1.1}$$

holds for all $x,y,z\in X$. Equivalently, there exists $c>0$ so that

$$\sum_{i=1}^N d(x_i,x_{i-1})^p\ge c\,d(x_0,x_N)^p\tag{1.2}$$

whenever $x_0,x_1,\dots,x_N$ is a finite chain of points in X. (For the equivalence of these definitions, see Section 2.) If (X,d) is a p-snowflake for some $p>1$ we say merely that (X,d) is a *snowflake*.

In this paper we give several geometric characterizations for snowflake spaces. These include a quantitative gap condition for roughly collinear points in the space, and a qualitative condition on the absence of rectifiable curves in all weak tangents (in the sense of Gromov) of the space. We also prove that certain classical self-similar planar sets are snowflakes.

2000 Mathematics Subject Classification: 54E35; Secondary 28A80, 30C65.

The first author was supported by the National Science Foundation under Award No. DMS-0228807. The second author was supported by the National Science Foundation under Award No. DMS-0070312.

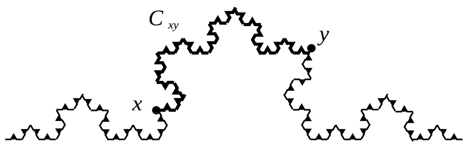

Figure 1. The von Koch snowflake curve.

Our terminology stems from the observation that the classical von Koch snowflake curve C, endowed with the planar Euclidean metric, is a p-snowflake with $p = \log 4/\log 3$. Indeed, we may choose $d'(x,y) = \mathscr{H}^p(C_{xy})^{1/p}$, where $\mathscr{H}^p$ denotes the Hausdorff p-measure on C and C_{xy} denotes the minimal connected subset of C containing x and y. See Figure 1.

A well-known theorem of Assouad [1] asserts that every doubling metric space may be realized as a snowflake subspace of some finite-dimensional Euclidean space. More precisely, for every doubling metric space (X,d) and every $0<\varepsilon<1$, the metric space (X,d^ε) admits a bi-Lipschitz embedding into some Euclidean space $\mathbf{R}^N$. The image of X in $\mathbf{R}^N$ is a p-snowflake with $p=1/\varepsilon$.

A theorem of Semmes (Theorem 6.3 in [11]) asserts that each p-snowflake set $A \subset \mathbf{R}^N$ generates a strong A_∞ weight

$$\omega(x) = \operatorname{dist}(x,A)^{N(p-1)}$$

on $\mathbf{R}^N$, and the distance function D_ω is bi-Lipschitz equivalent with the restriction of d_E^p to A. Here D_ω is the associated distance function defined by

$$D_\omega(x,y) := \inf \int_\gamma \omega^{1/N}\, ds,$$

the infimum being taken over all rectifiable curves γ in $\mathbf{R}^n$ joining x to y, and d_E denotes the Euclidean metric on $\mathbf{R}^N$. Combining this result of Semmes with the theorem of Assouad shows that every doubling metric space may be realized up to a bi-Lipschitz map within $(\mathbf{R}^N, D_\omega)$ for some metric D_ω associated with a strong A_∞ weight ω (Theorem 1.15 of [11]).

For the statement of our main result, we introduce the following separation-type condition. Let (X,d) be a metric space. For $x,y\in X$ and constants $0<\lambda<1$ and $\delta>0$, let

$$L(x,y;\lambda,\delta) := B\big(x,(\lambda+\delta)d(x,y)\big) \cap B\big(y,(1-\lambda+\delta)d(x,y)\big). \tag{1.3}$$

Thus $L(x,y;\lambda,\delta)$ is a "lens-shaped" set consisting of the intersection of two overlapping balls centered at x and y. Here and throughout the paper we denote by $B(x,r)$ the closed ball in X with center x and radius r. In case the underlying space X needs to be mentioned we write $L(x,y;\lambda,\delta) = L_X(x,y;\lambda,\delta)$.

Definition 1.4. We say that a metric space (X,d) is δ*-uniformly non-convex* (*UNC*), $0<\delta<\frac12$, if to each pair of points $x,y\in X$, there corresponds a value $\lambda=\lambda_{xy}\in(\delta,1-\delta)$ so that $L(x,y;\lambda,\delta)$ is empty. We say that (X,d) is *uniformly non-convex* if it is δ-uniformly non-convex for some $\delta>0$.

Theorem 1.5. *Let (X,d) be a compact, doubling metric space that admits a bi-Lipschitz embedding into a uniformly convex Banach space. Then the following are equivalent*:

(1) (X,d) *is a snowflake*;
(2) (X,d) *is uniformly non-convex*;
(3) *Each weak tangent of* (X,d) *contains no rectifiable curves.*

A pointed metric space $(Z_\infty, d_\infty, z_\infty)$ is called a *weak tangent* of (X,d) if there exist points $x_m \in X$ and positive reals r_m so that the pointed metric spaces $(X, r_m^{-1}d, x_m)$ converge in the topology of pointed Gromov–Hausdorff convergence to $(Z_\infty, d_\infty, z_\infty)$. We do not require that $r_m \to 0$; thus the original metric space (X,d) with suitable basepoint is itself a weak tangent. See [5, Chapters 8 and 9] or [4, Chapters 7 and 8] for further information on Gromov–Hausdorff convergence and weak tangent spaces.

A reformulation of Assouad's theorem states that every doubling snowflake space admits a bi-Lipschitz embedding in some finite-dimensional Euclidean space. It follows that if a compact, doubling metric space satisfies the hypotheses of Theorem 1.5 as well as either condition (2) or condition (3), then it admits a bi-Lipschitz embedding in a finite-dimensional Euclidean space. The question of which metric spaces admit a bi-Lipschitz embedding into a finite-dimensional Euclidean space is a well-known open problem in geometric analysis.

The implication (2) $\Rightarrow$ (1) in Theorem 1.5 holds for all metric spaces, and the implication (2) $\Rightarrow$ (3) holds for all separable spaces. In the absence of compactness and the doubling property, the implication (3) $\Rightarrow$ (2) fails for a certain bounded subset of l_2. These and other results of a similar nature appear in various sections of this paper. We do not know in what generality the implication (1) $\Rightarrow$ (2) holds.

As an application, we show that certain self-similar sets are snowflakes. We consider the *polygaskets* $\mathrm{PG}(N)$, $N \geq 3$. The set $\mathrm{PG}(N)$ is obtained as the invariant set for a collection of contractive similarities of a regular N-gon. For the precise definition, see Section 6. PG(3) is the standard *Sierpinski gasket*, while PG(4) is a closed square. Figure 2 shows $\mathrm{PG}(N)$ for $N = 3,5,6,8,9$.

Theorem 1.6. $\mathrm{PG}(N)$ *is a snowflake if and only if* $N = 5$ *or* $N \geq 7$.

The cases $N = 3,4,6$ are ruled out since the corresponding sets contain nontrivial line segments and hence violate condition (3) in Theorem 1.5.

To conclude, we mention a connection with conformal geometry. The *conformal dimension* of a metric space (X,d) is defined as the infimum of the Hausdorff dimensions of X with respect to all metrics quasisymmetrically (qs) equivalent with d:

$$\mathscr{C}\dim X = \inf\bigl\{\dim(X,d') : d \text{ and } d' \text{ are qs equivalent}\bigr\}.$$

For the definition and basic properties of quasisymmetric maps, see [6]. Since bi-Lipschitz equivalence implies quasisymmetric equivalence and the snowflaking operation $d \mapsto d' := d^\varepsilon$, $\varepsilon < 1$, generates metrics quasisymmetrically equivalent with d, $\mathscr{C}\dim X \le \dim(X,d)/p$ if (X,d) is a p-snowflake, and we may record the following.

Corollary 1.7. *Let (X,d) be a compact and doubling metric space which satisfies one of the equivalent conditions in Theorem* 1.5. *Then (X,d) is not minimal for conformal dimension, i.e.,* $\dim(X,d) > \mathscr{C}\dim X$.

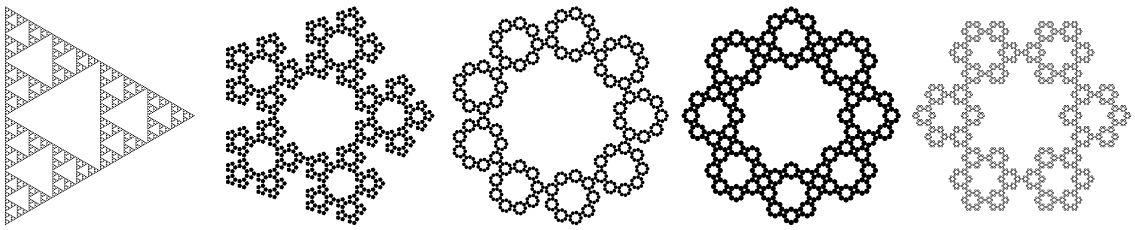

Figure 2. Polygaskets $\mathrm{PG}(N)$ for $N = 3,5,6,8,9$.

It follows from recent work of Keith and Laakso [7] that, under an additional regularity assumption, the inequality $\dim(X,d) > \mathscr{C}\dim X$ holds if and only if, for every $\alpha \ge 1$, the α-modulus of curve families in every weak tangent $(Z_\infty, d_\infty, z_\infty)$ to (X,d) is trivial. Now the nonexistence of rectifiable curves certainly implies the triviality of the α-modulus, so Corollary 1.7 follows from [7]. On the other hand, in the setting of Corollary 1.7 we have the strict inequalities

$$\dim(X,d) > \dim(X,d^p) > \mathscr{C}\dim X$$

for some $p > 1$.

The result of Keith and Laakso implies that $\mathrm{PG}(N)$ is not minimal for conformal dimension when $N \ge 3$, $N \ne 4$. Corollary 1.7 gives an alternate proof for this fact, apart from the cases $N = 3,6$. In [12] the authors proved that $\inf \dim f\big(\mathrm{PG}(N)\big) = \mathscr{C}\dim \mathrm{PG}(N) = 1$ if $N \not\equiv 0 \pmod 4$, where the infimum is taken over all quasiconformal maps $f\colon \mathbf{R}^2 \to \mathbf{R}^2$.

Organization of the paper. In Section 2 we prove that (1.2) is equivalent with the snowflake property. In Section 3 we prove the implication (2) $\Rightarrow$ (1) from Theorem 1.5 for general metric spaces, and discuss some related geometric conditions.

In Section 4 we specialize to uniformly convex Banach spaces and prove the equivalence of conditions (1) and (2).

In Section 5 we discuss Gromov–Hausdorff convergence and weak tangents, and complete the proof of Theorem 1.5.

In Section 6 we prove Theorem 1.6.

In an appendix we sketch the proof of a different geometric characterization of snowflake spaces due to Tomi Laakso. We thank him for suggesting that we include it here.

Notation. We remind the reader that we use the notation $B(x,r) = B_X(x,r)$ to denote the closed ball with center x and radius r in a metric space $X = (X,d)$. We write $\operatorname{diam} A$ for the diameter of a set $A \subset X$, and we write $\operatorname{dist}(A,B)$ for the distance between two sets $A, B \subset X$. In the case when $A = \{x\}$ is a singleton, we write $\operatorname{dist}(x,B) = \operatorname{dist}(\{x\},B)$. For $\varepsilon > 0$ and $A \subset X$, we write $N_\varepsilon(A) = \{x \in X : \operatorname{dist}(x,A) < \varepsilon\}$ for the ε-neighborhood of A.

We say that (X,d) is a *proper* metric space if each closed ball in X is compact, and that (X,d) is a *doubling space* if there exists a constant $M < \infty$ so that every ball of radius r in X may be covered by at most M balls of radius $\frac{1}{2}r$. Every complete doubling space is proper.

A map $f\colon (X,d) \to (X',d')$ is said to be L*-bi-Lipschitz*, $L \geq 1$, if the double inequality

$$d(x,y)/L \leq d'\bigl(f(x),f(y)\bigr) \leq Ld(x,y)$$

holds for all $x, y \in X$. If f is L-bi-Lipschitz for some $L < \infty$ we say that f is *bi-Lipschitz.* Two metrics d and d' on a given set X are called *bi-Lipschitz equivalent* if the identity map from (X,d) to (X,d') is bi-Lipschitz. 1-bi-Lipschitz maps are also called *isometries.*

For the purposes of this paper, a *curve* is a nonconstant continuous map γ from a compact interval $I = [a,b] \subset \mathbf{R}$ into X. We make the usual identification of the map γ and the image set $\gamma(I) \subset X$. A curve γ is *rectifiable* if its length

$$\operatorname{length}(\gamma) = \sup \sum_{i=1}^{N} d\bigl(\gamma(t_i, t_{i-1})\bigr)$$

is finite, where the supremum is taken over all ordered chains of points $a = t_0 < t_1 < \cdots < t_{N-1} < t_N = b$. For a subinterval $I_1 \subset I$ we write $\gamma|_{I_1}$ for the restriction of γ to I_1. If γ is an isometric embedding of I in X, the curve is a *geodesic.* If γ is an L-bi-Lipschitz embedding, the curve is an L*-quasigeodesic.*

For points v, w in a normed vector space $(V, \|\cdot\|)$, we write

$$[v,w] = \{tv + (1-t)w : 0 \leq t \leq 1\}$$

for the line segment joining v to w.

Finally, throughout this paper we denote by $C, C_1, c, c_1, \ldots$, various positive constants. The values of constants may change, even within a single line.

2. L^p-metrics and snowflake spaces

Let (X,d) be a metric space. We say that d is an L^p*-metric* if d satisfies the appropriate inequality from (1.1) for all $x, y, z \in X$. Note that if d is an L^p-metric, then d is an L^q-metric for all $1 \leq q \leq p \leq \infty$.

Recall from the introduction that (X,d) is a p-snowflake if d is bi-Lipschitz equivalent with an L^p-metric on X.

L^∞-metrics are more commonly known as *ultrametrics* (also *non-Archimedean* or *isosceles* metrics). David and Semmes [5, Proposition 15.7] have characterized those metrics on a space X which are bi-Lipschitz equivalent to an ultrametric. A metric space (X,d) is said to be *uniformly disconnected* if there exists $c>0$ so that

$$\max_{i=1,\dots,N} d(x_i,x_{i-1}) \ge c d(x,y) \tag{2.1}$$

for all finite chains of points $x=x_0,x_1,\dots,x_N=y$. Equivalently, there exists $c_1>0$ so that each ball $B(x,r)$ in X contains a set A satisfying $A\supset B(x,c_1 r)$ and $\operatorname{dist}(A,X\setminus A)\ge c_1 r$. (See [6, Exercise 14.26].)

Proposition 2.2 (David–Semmes). *Let (X,d) be a metric space. Then d is bi-Lipschitz equivalent with an ultrametric on X if and only if (X,d) is uniformly disconnected.*

Condition (1.2) is the L^p analogue of (2.1). The corresponding L^p version of Proposition 2.2 is the following.

Proposition 2.3. *Let $1\le p<\infty$. A metric space (X,d) is a p-snowflake if and only if there exists $c>0$ so that* (1.2) *holds for all finite chains $x_0,x_1,\dots,x_N$ in X.*

The proof of Proposition 2.3 is straightforward. Indeed, if d' is an L^p metric then

$$\sum_{i=1}^N d'(x_i,x_{i-1})^p \ge d'(x,y)^p$$

for all finite chains $x=x_0,x_1,\dots,x_N=y$. If (X,d) is a p-snowflake, then d is bi-Lipschitz equivalent with an L^p-metric and (1.2) follows. For the converse, assume that (1.2) holds for all finite chains. Set

$$d'(x,y) := \inf\left\{\sum_{i=1}^N d(x_i,x_{i-1})^p\right\}^{1/p},$$

the infimum being taken over all finite chains $x=x_0,x_1,\dots,x_N=y$. Then $d'\le d$ and it is immediate that d' satisfies the triangle inequality. By (1.2) $d'\ge c^{1/p}d$. Thus d' is an L^p-metric on X which is bi-Lipschitz equivalent with d.

3. Snowflake spaces and uniform non-convexity I

Theorem 3.1. *Let (X,d) be a metric space. If (X,d) is δ-uniformly non-convex for some $\delta>0$, then (X,d) is a p-snowflake with*

$$p = \frac{\log 2}{\log 2 - \log(1+4\delta^2)} > 1. \tag{3.2}$$

Observe that $p=p(\delta)$ in (3.2) tends to one as $\delta\to 0$, and tends to $+\infty$ as $\delta\to\frac12$.

Theorem 3.3. *Let (X,d) be a metric space. If (X,d) is a snowflake, then $f(X)$ is uniformly linearly non-convex in V whenever $f\colon X\to V$ is a bi-Lipschitz embedding of X into a normed vector space V.*

Definition 3.4. We say that a set A in a normed vector space $(V,\|\cdot\|)$ is η*-uniformly linearly non-convex* (*ULNC*) *in* V, $\eta>0$, if to each pair of points $a,b\in A$ there corresponds $c\in[a,b]$ so that $B_V(c,\eta\|a-b\|)$ is disjoint from A. If A is η-uniformly linearly non-convex in V for some $\eta>0$, we say that A is *uniformly linearly non-convex in* V.

We record the following corollary to Theorems 3.1 and 3.3.

Corollary 3.5. *Let (X,d) be a UNC space and $f\colon X\to V$ be a bi-Lipschitz embedding into a normed vector space V. Then $f(X)$ is ULNC in V.*

The converses of Theorem 3.1 and Corollary 3.5 fail to hold in general. See Example 4.6 for a compact and doubling snowflake subset of $(\mathbf{R}^2,\|\cdot\|_\infty)$ which is ULNC but not UNC. If X embeds bi-Lipschitzly in a uniformly convex Banach space, then the converses of Theorem 3.1 and Corollary 3.5 hold.

Proof of Theorem 3.1. Let (X,d) be δ-UNC for some $0<\delta<\frac{1}{2}$. We will show that (1.2) holds with p as in (3.2) and

$$c=a(\delta)^p,\qquad a(\delta):=4\delta\big(\tfrac{1}{2}-\delta\big). \tag{3.6}$$

We argue by induction on the length of the chain. If $N=1$ then (1.2) obviously holds for any $p\ge 1$ and any $c\le 1$, in particular, for the values in (3.2) and (3.6). Assume then that (1.2) holds for all chains of length at most $N-1$ and let $x_0,x_1,\dots,x_N$ be a chain of length N. Choose $\lambda=\lambda_{x_0x_N}\in(\delta,1-\delta)$ so that $L(x_0,x_N;\lambda,\delta)=\emptyset$. Since $L(x_0,x_N;\lambda,\delta)$ is empty, one of the following two alternatives must hold:

(i) For each $k\in\{0,\dots,N\}$, either

$$d(x_0,x_k)\le(\lambda-\delta+4\delta^2)d(x_0,x_N)$$

or

$$d(x_0,x_k)\ge(\lambda+\delta)d(x_0,x_N).$$

(ii) There exists an index $k\in\{1,\dots,N-1\}$ so that

$$d(x_0,x_k)+d(x_k,x_N)\ge(1+4\delta^2)d(x_0,x_N); \tag{3.7}$$

In case (i), there must exist $l\in\{1,\dots,N\}$ so that

$$d(x_0,x_{l-1})\le(\lambda-\delta+4\delta^2)d(x_0,x_N)$$

and

$$d(x_0,x_l)\geq(\lambda+\delta)d(x_0,x_N).$$

Then

$$\sum_{i=1}^{N} d(x_i,x_{i-1})^p \geq d(x_l,x_{l-1})^p \geq a(\delta)^p d(x_0,x_N)^p.$$

Suppose then that case (ii) holds. By the induction hypothesis,

$$\sum_{i=1}^{k} d(x_i,x_{i-1})^p \geq a(\delta)^p d(x_0,x_k)^p$$

and

$$\sum_{i=k+1}^{N} d(x_i,x_{i-1})^p \geq a(\delta)^p d(x_k,x_N)^p$$

whence

$$\sum_{i=1}^{N} d(x_i,x_{i-1})^p \geq a(\delta)^p \bigl(d(x_0,x_k)^p + d(x_k,x_N)^p\bigr).$$

By Lemma 3.8 and (3.7), the expression in parentheses on the right-hand side is at least $d(x_0,x_N)^p$. Again, (1.2) is satisfied. This completes the proof of the induction step and consequently the proof of Theorem 3.1. □

Lemma 3.8. *If* $\delta_1>0$ *and* $p=\log 2/\bigl(\log 2-\log(1+\delta_1)\bigr)$, *then* $A^p\leq B^p+C^p$ *whenever* $A,B,C\geq 0$ *satisfy* $(1+\delta_1)A\leq B+C$.

Proof. By homogeneity and a scaling argument it clearly suffices to consider the case $A=1$ and $B+C=1+\delta_1$. The minimum of $F_p(B):=B^p+(1+\delta_1)^p$ for $B\in[0,1+\delta_1]$ occurs at the midpoint of this interval, where $F_p\bigl(\frac12(1+\delta_1)\bigr)=2\bigl(\frac12(1+\delta_1)\bigr)^p=1$. □

Proof of Theorem 3.3. Let (X,d) be a p-snowflake for some $p>1$. By Proposition 2.3, there exists $c>0$ so that (1.2) holds for all finite chains of points. For fixed $L<\infty$, choose an integer N so that

$$N\Bigl(\frac{2L^2}{N}\Bigr)^p<c,$$

and set $\eta=1/(2N)$. Let $f\colon X\to V$ be an L-bi-Lipschitz embedding of X into a normed vector space V; we will show that $f(X)$ is η-ULNC.

Let $x,y\in X$ and choose a chain of equally spaced collinear points

$$f(x)=z_0,z_1,\dots,z_N=f(y)$$

in the ambient space V. For $i=0,\dots,N$, set $B_i := B_V\bigl(z_i, \|f(x)-f(y)\|/2N\bigr)$. Suppose that each ball $B_1,\dots,B_{N-1}$ meets $f(X)$, and choose $f(x_i)\in B_i\cap f(X)$, $x_i\in X$, $i=0,\dots,N$. Set $x_0=x$ and $x_N=y$. For each i we have

$$\begin{aligned}\|f(x_i)-f(x_{i-1})\| &\le \|f(x_i)-z_i\| + \|z_i-z_{i-1}\| + \|z_{i-1}-f(x_{i-1})\|\\ &\le \frac{2}{N}\|f(x_0)-f(x_N)\|\end{aligned}$$

and so

$$\|x_i-x_{i-1}\| \le \frac{2L^2}{N}\|x_0-x_N\|.$$

Thus

$$\sum_{i=1}^{N}\|x_i-x_{i-1}\|^p \le N\biggl(\frac{2L^2}{N}\biggr)^p\|x_0-x_N\|^p < c\|x_0-x_N\|^p$$

which contradicts (1.2). It follows that one of the balls $B\bigl(z_k, \|f(x)-f(y)\|/(2N)\bigr)$ is disjoint from $f(X)$, which demonstrates that $f(X)$ is η-ULNC in V. □

4. Snowflake spaces and uniform non-convexity II

Recall that a normed vector space $(V,\|\cdot\|)$ is called *uniformly convex* if to each $0<\varepsilon\le 2$ there corresponds $\delta(\varepsilon)>0$ so that $\|v-w\|<\varepsilon$ whenever $\|v\|=\|w\|=1$ and $\bigl\|\frac12(v+w)\bigr\|>1-\delta(\varepsilon)$. Laakso proved that uniform convexity is equivalent with the so-called round ball condition [9].

Definition 4.1. A metric space (X,d) is called a *strong round ball* (*SRB*) *space* if for every $\varepsilon>0$ there exists $\delta_1(\varepsilon)>0$ so that for all $x,y\in X$ and all $0<\lambda<1$,

$$\operatorname{diam} L\bigl(x,y;\lambda,\delta_1(\varepsilon)\bigr) \le \varepsilon d(x,y), \tag{4.2}$$

where $L(x,y;\lambda,\delta)$ is the lens set introduced in (1.3).

If (4.2) holds only for $\lambda=\frac12$, we say that (X,d) is a *round ball* (*RB*) *space.* This is the original notion due to Laakso, see Definition 1.0 of [9]. In the setting of Banach spaces, SRB and RB are equivalent:

Proposition 4.3. *For a Banach space $(V,\|\cdot\|)$, the following conditions are equivalent*:

(i) *V is uniformly convex*;
(ii) *$(V,\|\cdot\|)$ is an RB space*;
(iii) *$(V,\|\cdot\|)$ is an SRB space*;
(iv) *to each $\varepsilon>0$ there corresponds $\delta_2(\varepsilon)>0$ so that $\|q-\|q\|\cdot p\|<\varepsilon$ whenever $p,q\in V$ satisfy $\|p\|=1$ and $\|q\|+\|p-q\|<1+\delta_2(\varepsilon)$.*

Before giving the proof of this proposition, we show its connection with the characterization question for snowflake spaces.

Proposition 4.4. *Let* $(V,\|\cdot\|)$ *be a uniformly convex Banach space. Then* A *is ULNC in* V *if and only if* A *is UNC.*

Proof of Proposition 4.4. Let $A\subset V$ be UNC. For any $a,b\in A$, there is $c=(1-\lambda)a+\lambda b$, $0<\lambda<1$, so that

$$B_V(c,\eta\|a-b\|)\cap A=\emptyset.$$

By Proposition 4.3 $(V,\|\cdot\|)$ is an SRB space, whence

$$\operatorname{diam} L_V\bigl(a,b;\lambda,\delta_1(\eta)\bigr)\le\eta\|a-b\|.$$

Thus

$$L_V\bigl(a,b;\lambda,\delta_1(\eta)\bigr)\subset B_V(c,\eta\|a-b\|)$$

and so

$$L_A\bigl(a,b;\lambda,\delta_1(\eta)\bigr)=L_V\bigl(a,b;\lambda,\delta_1(\eta)\bigr)\cap A=\emptyset.$$

This proves that A is $\delta_1(\eta)$-ULNC in V. The converse assertion follows from Corollary 3.5. □

As a corollary, we obtain the equivalence of conditions (1) and (2) in Theorem 1.5.

Corollary 4.5. *Let* (X,d) *be a metric space which admits a bi-Lipschitz embedding into a uniformly convex Banach space* $(V,\|\cdot\|)$. *Then* X *is a snowflake space if and only if the image of* X *in* V *is ULNC.*

Proof of Corollary 4.5. The "only if" statement is Theorem 3.3. For the converse, suppose that the image of X in V is ULNC. By Proposition 4.4, $f(X)$ is a UNC metric space, with metric induced from V. Theorem 3.1 ensures that $f(X)$ is a snowflake. The conclusion follows since the snowflake condition is bi-Lipschitz invariant. □

Uniform non-convexity and uniform linear non-convexity are not bi-Lipschitz invariant conditions. For subspaces of uniformly convex Banach spaces, uniform non-convexity is equivalent with the snowflake condition. This raises the question: can every snowflake space be isometrically embedded in some uniformly convex Banach space? Every ultrametric space may be isometrically embedded in a Hilbert space; see Corollary 1.3 in [10] and compare with Problem 3 from [10].

4.6. The von Koch snowflake is not a UNC subspace of $(\mathbf{R}^2, |\cdot|_\infty)$. Let Λ be the standard von Koch snowflake curve in $\mathbf{R}^2$ (see Figure 1). We assume that the endpoints of Λ are at the origin and $e_1 = (1,0)$, and that Λ is contained in the triangle $T = \{(x,y) : 0 \le y \le \min\{|x|, |1-x|\}\}$.

Let d be the standard Euclidean metric on $\mathbf{R}^2$, and let $|\cdot|_\infty$ be the maximum metric on $\mathbf{R}^2$. As indicated in the introduction, (Λ, d) is a snowflake; since $|\cdot|_\infty$ and d are bi-Lipschitz equivalent, $(\Lambda, |\cdot|_\infty)$ is also a snowflake, and hence is ULNC in $(\mathbf{R}^2, |\cdot|_\infty)$.

We claim that $(\Lambda, |\cdot|_\infty)$ is not UNC. To see this, it suffices to note that the lens sets $L(0, e_1, \lambda, \delta)$ are nonempty for every $0 \le \lambda \le 1$ and $\delta > 0$. In fact, for every such λ there exists $z \in \Lambda$ with $|z|_\infty = \lambda$ and $|z - e_1|_\infty = 1 - \lambda$. Indeed, since Λ is a closed curve joining the origin to e_1, there exists $z = (x,y) \in \Lambda$ with $|z|_\infty = \lambda$. Since $\Lambda \subset T$, we must have $x = \lambda$, in which case $|z - e_1|_\infty = 1 - x = 1 - \lambda$.

Proof of Proposition 4.3. (i) $\Leftrightarrow$ (ii): This is Lemma 5.2 from [9].

(iii) $\Rightarrow$ (ii): This is trivial.

(i) $\Rightarrow$ (iii): We will show that the SRB condition holds with

$$\delta_1 = \delta_1(\varepsilon) := \frac{1}{25} \min\{1, \varepsilon^2, \delta(\varepsilon)^2\}, \tag{4.7}$$

where $\delta(\varepsilon)$ is the function in the definition of uniform convexity.

Let $x, y \in V$ and $0 < \lambda < 1$, and let $a, b \in L := L(x, y; 1-\lambda, \delta_1)$. Our goal is to show that $\|a - b\| \le \varepsilon \|x - y\|$. Since the SRB condition is scale-invariant, we may assume that $x = 0$ and $\|y\| = 1$. Moreover, we may assume that $\lambda \ge \sqrt{\delta_1} - \delta_1$; otherwise $\|a - b\| \le 2(\lambda + \delta_1) \le 2\sqrt{\delta_1} \le \varepsilon$.

By the definition of L, we have $\|a\|, \|b\| \le 1 - \lambda + \delta_1$ and $\|a - y\|, \|b - y\| \le \lambda + \delta_1$. Since L is convex, $\frac{1}{2}(a+b) \in L$ whence

$$\left\|\tfrac{1}{2}(a+b)\right\| \ge \lambda - \delta_1. \tag{4.8}$$

Suppose that $\|a - b\| > \varepsilon$. Then $v = a/\|a\|$ and $w = b/\|b\|$ are elements of the unit sphere of V, and a calculation shows

$$\|v - w\| \ge \frac{\varepsilon}{\lambda + \delta_1} - \frac{2\delta_1}{\lambda - \delta_1} \ge \varepsilon.$$

By the uniform convexity of V, $\left\|\frac{1}{2}(v+w)\right\| \le 1 - \delta(\varepsilon)$. Writing

$$v + w = \frac{1}{\|a\|}(a+b) + \frac{\|a\| - \|b\|}{\|a\|\,\|b\|} b$$

and using the definition of v and w together with the bound in (4.8), we derive the inequality

$$\frac{\lambda - \delta_1}{\lambda + \delta_1} \le 1 - \delta(\varepsilon) + \frac{\delta_1}{\lambda - \delta_1},$$

which can not be true due to the choice of δ_1 in (4.7). This completes the proof of the implication (iii) $\Rightarrow$ (i).

(iii) $\Rightarrow$ (iv): We will show that the condition in (iv) holds with

$$\delta_2(\varepsilon) := \min\bigl\{1, \tfrac{1}{2}\varepsilon, \delta_1\bigl(\tfrac{1}{2}\varepsilon\bigr)\bigr\}.$$

Suppose that $p, q \in V$ with $\|p\| = 1$ and $\|q\| + \|p - q\| < 1 + \delta_2$. Applying the SRB condition with $x = 0$, $y = p$ and $\lambda = \|q\|/(1 + \delta_2)$, we conclude that

$$\operatorname{diam} L(0, p; \lambda, \delta_2) \le \tfrac{1}{2}\varepsilon.$$

Since λp and q belong to this lens set, we find that $\|q - \lambda p\| \le \frac{1}{2}\varepsilon$, whence

$$\|q - \|q\| \cdot p\| \le \|q - \lambda p\| + \bigl|\lambda - \|q\|\bigr| \le \frac{\varepsilon}{2} + \frac{\delta_2}{1 + \delta_2}\|q\| < \frac{\varepsilon}{2} + \delta_2 \le \varepsilon$$

as desired.

(iv) $\Rightarrow$ (iii): We will show that the strong round ball condition holds with

$$\delta_1 = \delta_1(\varepsilon) := \tfrac{1}{4}\min\bigl\{1, \varepsilon, \delta_2\bigl(\tfrac{1}{4}\varepsilon\bigr)\bigr\}, \tag{4.9}$$

where $\delta_2(\varepsilon)$ is the function in condition (iv).

Let $x, y \in V$ and $0 < \lambda < 1$. As before, we may assume without loss of generality that $x = 0$ and $\|y\| = 1$. Let $p = y$ and let $q \in L(0, p, \lambda, \delta_1)$. Then

$$\|q\| + \|p - q\| < 1 + 2\delta_1 < 1 + \delta_2\bigl(\tfrac{1}{4}\varepsilon\bigr).$$

By condition (iv), $\bigl\|q - \|q\|p\bigr\| < \frac{1}{4}\varepsilon$ and so

$$\operatorname{diam} L(0, p, \lambda, \delta_1) \le 2\bigl(\tfrac{1}{4}\varepsilon + \delta_1\bigr) \le \varepsilon$$

as desired. □

5. Snowflake spaces and weak tangents

Definition 5.1. Let (X_m, d_m, p_m) be a sequence of pointed metric spaces. We say that (X_m, d_m, p_m) *converges* to a limit space (X, d, p) (in the sense of pointed Gromov–Hausdorff convergence) if for each $R > 0$ and $\varepsilon > 0$ there exists $M = M(\varepsilon, R)$ and maps $f_m = f_m^\varepsilon \colon B(p_m, R) \to X$, $m \ge M$, satisfying the following three conditions:

(i) $f_m(p_m) = p$;
(ii) $\bigl|d\bigl(f_m(x), f_m(x')\bigr) - d_m(x, x')\bigr| < \varepsilon$ for all $x, x' \in B(p_m, R)$;
(iii) $B(p, R - \varepsilon) \subset N_\varepsilon\bigl(f_m\bigl(B(p_m, R)\bigr)\bigr)$.

We denote this form of convergence by $(X_m, p_m) \xrightarrow{GH} (X, p)$.

This is a natural formulation of Gromov–Hausdorff convergence for (possibly) unbounded spaces. See, for example, Chapter 8 of [4]. In case the spaces X_m, X have uniformly bounded diameters, it agrees with the usual definition [4, Exercise 8.1.2]. If the metric spaces X_m are proper and the limit space X is complete, then X is proper and uniquely determined (up to pointed isometry).

In the technical language which has become standard in this subject [2], [3], condition (ii) states that f_m is an ε-rough isometric embedding of $B(p_m, r)$ into X.

Definition 5.2. We say that a pointed metric space $(Z_\infty, d_\infty, z_\infty)$ is a *weak tangent* of a metric space (X, d) if there exist sequences $p_m \in X$ and $0 < r_m \le C < \infty$ so that $(X, r_m^{-1} d, p_m) \xrightarrow{GH} (Z_\infty, d_\infty, z_\infty)$.

See, for example, Chapter 9 of [5].

If the scaling ratios r_m in Definition 5.2 are uniformly bounded away from zero, then the limit space (Z_∞, d_∞) is isometric with $(X, c^{-1}d)$ for some suitable $c \in (0, \infty)$. This is easy to prove.

Proposition 5.3. *Let (X, d) be a metric space which admits an L-bi-Lipschitz embedding in a Banach space $(V, \|\cdot\|)$. Then for each weak tangent $(Z_\infty, d_\infty, p_\infty)$ of (X, d) and each $R > 0$, the ball $B(z_\infty, R)$ in Z_∞ admits an L-bi-Lipschitz embedding in $(V, \|\cdot\|)$.*

Proof. We may assume that $(X, r_m^{-1} d, p_m) \xrightarrow{GH} (Z_\infty, d_\infty, z_\infty)$ for some $p_m \in X$ and $r_m \to 0$. Fix an L-bi-Lipschitz map $\varphi\colon X \to V$, and define a sequence of L-bi-Lipschitz embeddings φ_m of X into V by

$$\varphi_m(x) = r_m^{-1}\big(\varphi(x) - \varphi(p_m)\big).$$

Fix $R > 0$ and $0 < \varepsilon < R$. By Definition 5.1, for sufficiently large m, there exists $f_m^\varepsilon\colon B(p_m, 2r_m R) \to Z_\infty$ with $f_m^\varepsilon(p_m) = z_\infty$,

$$\left| d_\infty\big(f_m^\varepsilon(x), f_m^\varepsilon(x')\big) - \frac{d(x, x')}{r_m} \right| < \varepsilon, \tag{5.4}$$

and $B(z_\infty, R) \subset N_\varepsilon\big(f_m^\varepsilon\big(B(p_m, 2r_m R)\big)\big)$.

We now define a map h^ε from $B(z_\infty, R)$ into $V \times \mathbf{R}$ as follows. Fix a sufficiently large integer $m = m(\varepsilon)$ so that the map f_m^ε exists. Given $z \in B(z_\infty, R) \subset Z_\infty$, choose $x \in B(p_m, 2r_m R)$ with $d_\infty\big(z, f_m^\varepsilon(x)\big) < \varepsilon$. Define $h^\varepsilon(z)$ to be the point $\varphi_m(x) \in V$. From (5.4) we deduce that h^ε is an $(L, 3\varepsilon)$-rough quasi-isometric embedding of $B(z_\infty, R)$ into V, i.e.,

$$\frac{1}{L}\big(d_\infty(z, w) - 3\varepsilon\big) \le \|h^\varepsilon(z) - h^\varepsilon(w)\| \le L(d_\infty(z, w) + 3\varepsilon)$$

for all $z, w \in B(z_\infty, R)$. Applying a standard Cantor diagonal argument and passing to a subsequence if necessary, we construct an L-bi-Lipschitz embedding $h\colon B(z_\infty, R) \to V$. □

Remark 5.5. We do not know whether the full weak tangent Z_∞ may be embedded in V. Note that it is not guaranteed in the above proposition that the various embeddings of the balls $B(z_\infty, R)$ into V are coherent.

We now state the principal results of this section.

Theorem 5.6. *If (X,d) is a snowflake, then every weak tangent of (X,d) contains no rectifiable curves.*

Theorem 5.7. *Let (X,d) be a compact, doubling metric space that admits an L-bi-Lipschitz embedding into a Banach space $(V, \|\cdot\|)$ for some $L \ge 1$. If no weak tangent of (X,d) contains an L-quasigeodesic, then the image of X is ULNC in V.*

Proposition 5.8. *Let (X,d) be a complete doubling metric space which admits an L-bi-Lipschitz embedding into a uniformly convex Banach space. If some weak tangent of (X,d) contains a rectifiable curve, then some weak tangent of (X,d) contains an L-quasigeodesic.*

In Theorem 5.7 uniform convexity for V is not needed, while compactness and the doubling condition are used to apply Gromov's compactness theorem. Theorem 5.7 does not hold in the absence of the doubling and compactness assumption. See Example 5.13.

Theorems 5.6 and 5.7 complete the proof of Theorem 1.5. Indeed, the implication (1) $\Rightarrow$ (3) follows from Theorem 5.6, while the implication (3) $\Rightarrow$ (2) follows from Theorem 5.7.

In case (X,d) is a compact and doubling space, Proposition 5.8 is a simple corollary of Theorem 5.7, Theorem 4.5 and Theorem 5.6.

Proof of Theorem 5.6. The proof is very similar to the proof of Theorem 3.3. Assume that (X,d) is a p-snowflake for some $p > 1$, and suppose that some weak tangent $(Z_\infty, d_\infty, z_\infty)$ of (X,d) contains a rectifiable curve γ with endpoints a and b. Let $d = d_\infty(a,b)$, let l be the length of γ, and let c be the constant from (1.2). Choose an integer N so that

$$2N\left(\frac{l}{dN}\right)^p < c.$$

Let $a = z_0, z_1, \dots, z_N = b$ be an ordered sequence of points on γ satisfying

$$d_\infty(z_i, z_{i-1}) = \frac{l}{N}$$

for all $i = 1, \dots, N$.

Choose $p_m \in X$ and $r_m > 0$ so that $(X, r_m^{-1} d, p_m) \xrightarrow{GH} (Z_\infty, d_\infty, z_\infty)$. Applying condition (ii) of Definition 5.1 and passing to a subsequence if necessary, we find points

$$a_m := x_0^m, x_1^m, \dots, x_N^m =: b_m$$

in X so that

$$|r_m^{-1}d(x_i^m,x_j^m)-d_\infty(z_i,z_j)|<\frac{1}{m} \tag{5.9}$$

for each $0\le i<j\le N$. Thus

$$d(x_i^m,x_j^m)<r_m\biggl(\frac{|i-j|}{N}l+\frac{1}{m}\biggr)$$

for each $0\le i<j\le N$ and

$$cd(a_m,b_m)^p\le\sum_{i=1}^N d(x_i^m,x_{i-1}^m)^p<Nr_m^p\biggl(\frac{l}{N}+\frac{1}{m}\biggr)^p$$

by (1.2). From (5.9) we see that $d(a_m,b_m)>r_m\bigl(d-(1/m)\bigr)$ whence

$$c\biggl(d-\frac{1}{m}\biggr)^p<N\biggl(\frac{l}{N}+\frac{1}{m}\biggr)^p.$$

We obtain a contradiction upon passing to the limit as $m\to\infty$. □

Proof of Theorem 5.7. Let φ be a bi-Lipschitz embedding of X in Y. Since no weak tangent of (X,d) contains an L-quasigeodesic, no weak tangent of $(\varphi(X),\|\cdot\|)$ contains a geodesic. We use here the fact that L-bi-Lipschitz maps pass to weak tangents with no increase in the bi-Lipschitz constant. It thus suffices to prove the stated result under the assumptions that (X,d) is contained in V and no weak tangent of (X,d) contains a geodesic.

Suppose that (X,d) fails to be η-ULNC in V for any $\eta>0$. Then for each $m\in\mathbf{N}$ there exist points $x_m,y_m\in X$ so that the line segment $[x_m,y_m]\subset V$ is contained within a suitable neighborhood of X in V:

$$[x_m,y_m]\subset\biggl\{v\in V:\operatorname{dist}(v,X)<\frac{1}{m}d(x_m,y_m)\biggr\}. \tag{5.10}$$

Passing to a subsequence if necessary, we may assume that the sequences (x_m), (y_m) converge to points $x,y\in X$ respectively. We distinguish two cases:

Case I ($x\ne y$): The geodesic segments $[x_m,y_m]$ converge to $[x,y]$ in the Hausdorff metric on V. Combining this observation with (5.10) and using the fact that X is a closed subset of V, we conclude that $[x,y]\subset X$. Thus the original metric space (X,d) (which occurs as a weak tangent of (X,d); see the remark following Definition 5.2) contains a geodesic segment.

Case II ($x = y$): Set $r_m := d(x_m, y_m)$. By Gromov's compactness theorem [4, Theorem 8.1.10], the rescaled metric spaces $(X, r_m^{-1} d, x_m)$ converge after passing to an appropriate subsequence to a weak tangent $(Z_\infty, d_\infty, z_\infty)$. Applying Definition 5.1 with $R = 2$ and again passing to a subsequence if necessary, we deduce the existence of maps $f_m\colon B(x_m, 2r_m) \to Z_\infty$ satisfying $f_m(x_m) = z_\infty$ and

$$\left| d_\infty\big(f_m(x), f_m(x')\big) - \frac{d(x,x')}{r_m} \right| < \frac{1}{m}.$$

Set $w_m = f_m(y_m)$. Observe that

$$w_m \in B(z_\infty, 1 + m^{-1}) \setminus B(z_\infty, 1 - m^{-1}). \tag{5.11}$$

For each $\lambda \in [0,1]$ consider the point $v_{m,\lambda} := \lambda y_m + (1-\lambda) x_m \in V$. By (5.10) there exists $u_{m,\lambda} \in X$ with $\|u_{m,\lambda} - v_{m,\lambda}\| < r_m/m$. Then $u_{m,\lambda} \in B(x_m, 2r_m)$ and the quantity

$$f_m(u_{m,\lambda})$$

is well-defined as an element of Z_∞. Define a map from $[0,1]$ to Z_∞ by

$$\lambda \mapsto f_m(u_{m,\lambda}). \tag{5.12}$$

From the aforementioned properties of f_m, we deduce that

$$\big| d_\infty\big(f_m(u_{m,\lambda}), f_m(u_{m,\lambda'})\big) - |\lambda - \lambda'| \big| < \frac{3}{m},$$

i.e., the map in (5.12) defines a $(1, 3/m)$-rough quasi-isometric embedding of $[0,1]$ into Z_∞ sending 0 to z_∞ and 1 to w_m. A Cantor diagonal argument as in the proof of Proposition 5.3 yields an isometric embedding of a dense subset of $[0,1]$ into Z_∞ sending 0 to z_∞ and 1 to a suitable cluster point w_∞ of the sequence (w_m). This embedding may be extended to an isometric embedding of $[0,1]$ into Z_∞. Thus Z_∞ contains a geodesic segment. □

Proof of Proposition 5.8. We assume that (X, d) is a complete and doubling (hence proper) space which admits an L-bi-Lipschitz embedding into a uniformly convex Banach space $(V, \|\cdot\|)$, and that some weak tangent $(Z_\infty, d_\infty, z_\infty)$ of (X, d) contains a rectifiable curve γ. Since properness descends to weak tangents, we may include γ within a compact ball B. By Proposition 5.3, B admits an L-bi-Lipschitz embedding in V. Since B contains a rectifiable curve, Theorem 5.6 implies that (B, d_∞) is not a snowflake. Recall that the space itself occurs as a weak tangent. Since the doubling condition passes to weak tangents, we conclude from Theorem 5.7 that some weak tangent $(\widetilde{W}_\infty, \tilde\delta_\infty, \widetilde{w}_\infty)$ of (B, d_∞) contains an L-quasigeodesic $\tilde\gamma$. Suppose that $(\widetilde{W}_B, \tilde\delta_\infty, \widetilde{w}_\infty)$ of (B, d_∞) arises as the

Gromov–Hausdorff limit of a sequence $(B, r_i^{-1} d_\infty, z_i)$, $z_i \in B$. By Gromov's compactness theorem, a suitable subsequence of $(Z_\infty, r_i^{-1} d_\infty, z_i)$ converges to a weak tangent $(W_\infty, \delta_\infty, w_\infty)$ of (Z_∞, d_∞). Moreover, $(W_\infty, \delta_\infty)$ contains $(\widetilde{W}_\infty, \tilde{\delta}_\infty)$, and hence contains $\tilde{\gamma}$. Since weak tangents of weak tangents are weak tangents, $(W_\infty, \delta_\infty, w_\infty)$ is a weak tangent of the original space (X, d). This completes the proof of Proposition 5.8. □

5.13. A non-snowflake subspace of l^2 having no rectifiable curves in its weak tangents. For each $n \in \mathbf{N}$, let $2^{-n}\mathbf{Z}$ denote the collection of real numbers of the form $j \cdot 2^{-n}$, $j \in \mathbf{Z}$. Denote by Q_n the closed cube in $\mathbf{R}^n$ centered at $(2, 0, \ldots, 0)$ of side length $2n^{-1/2}$ with edges parallel to the coordinate axes, and let

$$X_n = Q_n \cap (2^{-n}\mathbf{Z})^n \subset \mathbf{R}^n.$$

Denote by i_n the embedding of $\mathbf{R}^n$ into l^2 given by

$$(x_1, \ldots, x_n) \mapsto \sum_{k=1}^{n} x_k e\left(\frac{n(n-1)}{2} + k\right),$$

where $\{e(1), e(2), \ldots\}$ denotes an orthonormal basis for l^2. Note that the images $i_n(\mathbf{R}^n)$ lie in orthogonal subspaces of l^2, and that $\operatorname{dist}\bigl(i_n(X_n), i_{n'}(X_{n'})\bigr) \geq \sqrt{2}$ for all $n \neq n'$.

Let $X = \bigcup_{n \in \mathbf{N}} i_n(X_n)$. We endow X with the metric d induced from l^2, and observe that X is a countable subset of a closed ball of radius 3. Moreover, X is neither compact nor doubling, since it contains an infinite set of points $i_1(2), i_2(2,0), i_3(2,0,0), \ldots$ with mutual distance $2\sqrt{2}$.

It is clear that X is not a snowflake space, since it contains collections of equally spaced collinear points of arbitrarily large cardinality. We claim however no weak tangent of X contains a rectifiable curve. In fact, every weak tangent of X is either isometric with a rescaled copy of X, or is a singleton.

Suppose that some sequence $(X, r_m^{-1} d, p_m)$ Gromov–Hausdorff converges to a limit space $(Z_\infty, d_\infty, z_\infty)$. We distinguish three cases: (i) $r_m \geq c > 0$, (ii) $r_m \to 0$ and $p_m \in \bigcup_{n=1}^{N} i_n(X_n)$ for some $N < \infty$ and all m, (iii) $r_m \to 0$ and $p_m \in i_{n(m)}(X_{n(m)})$ with $n(m) \to \infty$ as $m \to \infty$. Modulo restriction to a subsequence, one of these possibilities must occur.

Case (i) has already been mentioned (after Definition 5.2); the limit space (Z_∞, d_∞) must be isometric with a rescaled copy of (X, d). In case (ii) we claim that Z_∞ must be a singleton. This is a consequence of the following

Lemma 5.14. *Suppose that $(X, r_m^{-1} d, p_m)$ Gromov–Hausdorff converges to $(Z_\infty, d_\infty, z_\infty)$. Assume that $r_m \to 0$ and $\inf_m \operatorname{dist}(p_m, X \setminus \{p_m\}) > 0$. Then $Z_\infty = \{z_\infty\}$.*

We are left with case (iii): $r_m \to 0$ and $p_m \in i_{n(m)}(X_{n(m)})$ with $n(m) \to \infty$ as $m \to \infty$. Here we distinguish two further subcases: (a) $r_m/2^{-n(m)} \to 0$ and (b) $\limsup r_m/2^{-n(m)} > 0$. In the former case, we may use a scaled version of Lemma 5.14 to conclude that Z_∞ is a singleton. Observe that $\operatorname{dist}(p, X \setminus \{p\}) \geq 2^{-n}$ if $p \in i_n(X_n)$.

We claim that case (iii)(b) cannot occur; in essence, the sequence $(X, r_m^{-1}d, p_m)$ cannot be Gromov–Hausdorff Cauchy. Passing to a subsequence, we may assume that

$$r_m \geq \delta \cdot 2^{-n(m)}$$

for all m and some $0 < \delta < 1$. Choose $R = 16/\delta$ and $\varepsilon = 1/\delta$. For sufficiently large m, there exists a map f_m from the ball $B(p_m, r_m R)$ in $(X, \|\cdot\|_2)$ to Z_∞ with $f_m(p_m) = z_\infty$,

$$\left| d_\infty\big(f_m(x), f_m(x')\big) - \frac{\|x - x'\|_2}{r_m} \right| < \varepsilon, \tag{5.15}$$

and

$$B(z_\infty, R - \varepsilon) \subset N_\varepsilon f_m\big(B(p_m, r_m R)\big), \tag{5.16}$$

where $B(p_m, r_m R)$ denotes the ball in X centered at p_m with radius $r_m R$.

Assume that m, m' are sufficiently large that $\max\{r_m, r_{m'}\} < 1/R$. Then

$$B(p_m, r_m R) = B_{l^2}(p_m, r_m R) \cap X \subset i_{n(m)}(X_{n(m)})$$

and

$$B(p_{m'}, r_{m'} R) \subset i_{n(m')}(X_{n(m')}).$$

Note that

$$f_m\big(B\big(p_m, \tfrac{1}{2} r_m R\big) \subset B\big(z_\infty, \tfrac{1}{2}R + \varepsilon\big) \subset B(z_\infty, R - \varepsilon) \subset N_\varepsilon f_{m'}(B(p_{m'}, r_{m'} R))$$

by (5.15).

From the choice of the sets X_n, each of the balls $B\big(p_m, \frac{1}{2} r_m R\big) \subset X$ contains a set of cardinality at least $n(m)$, whose mutual distances are all at least $\frac{1}{2} r_m R$. Then (5.15) implies that $B(z_\infty, R - \varepsilon) \supset f_m\big(B\big(p_m, \frac{1}{2} r_m R\big)\big)$ contains a set of the same cardinality whose mutual distances are all at least $\frac{1}{2}R - \varepsilon$; (5.16) then implies that $f_{m'}\big(B\big(p_{m'}, r_{m'} R\big)\big)$ contains a set of the same cardinality whose mutual distances are all at least $\frac{1}{2}R - 3\varepsilon$. Finally, a second application of (5.15) shows that $B(p_{m'}, r_{m'} R)$ contains a set of $n(m)$ points whose mutual distances are all at least

$$r_{m'}\big(\tfrac{1}{2}R - 4\varepsilon\big) = \tfrac{1}{4} r_{m'} R.$$

Since $B(p_{m'}, r_{m'} R) \subset i_{n(m')}(X_{n(m')})$ which contains at most $2^{n(m')^2 + 2n(m') - 1}$ points, we obtain a contradiction provided m is sufficiently large. We conclude that case (iii)(b) cannot occur.

Remark 5.17. Modify the above construction by replacing each point z in $i_n(X_n)$ by a scaled von Koch snowflake arc of diameter 2^{-n-10} containing z. The resulting space Y is again not a snowflake space, and each weak tangent of Y is either a scaled copy of Y or a scaled von Koch snowflake arc.

6. The polygasket $\mathrm{PG}(N)$ is a snowflake if $N = 5$ or $N \geq 7$

The *polygasket* $\mathrm{PG}(N)$ is the self-similar subset of $\mathbf{R}^2$ characterized by the following iterative procedure: fix an initial regular N-gon $P_\emptyset$ with sides of unit length, called the *level* 0 *polygon*, and at the mth step, replace each level m polygon Π with N regular N-gons $\Pi_1, \ldots, \Pi_N$, pairwise congruent and contained in Π, which satisfy two properties:

(i) Each Π_i shares a single vertex with Π.
(ii) Π_i and Π_j intersect if and only if they meet a common side of Π. Label the Π_i's so that Π_i and Π_j are disjoint if and only if $|i-j| \geq 2 \pmod N$. If $|i-j| = 1 \pmod N$ and $N \neq 0 \pmod 4$ (the *post-critically finite* case), the intersection of Π_i and Π_j consists of a single point which is a vertex of both Π_i and Π_j. If $|i-j| = 1 \pmod N$ and $N = 0 \pmod 4$, the intersection of Π_i and Π_j consists of a line segment which is a common edge of both Π_i and Π_j.

We call the subpolygons $\Pi_1, \ldots, \Pi_N$ the *children* of Π. We denote the children of the level 0 polygon $P_\emptyset$ by $P_1, \ldots, P_N$. Each child Π_i of a given polygon Π has side length equal to r_N times the side length of Π. The scaling ratio r_N can be explicitly determined as a function of N; see Section 6 of [12].

For each m, let K_m denote the union of all level m polygons. Then

$$\mathrm{PG}(N) = \bigcap_{m=0}^{\infty} K_m.$$

Figure 2 shows $\mathrm{PG}(N)$ for $N = 3, 5, 6, 8, 9$.

In this subsection, we prove that $\mathrm{PG}(N)$ is a snowflake for $N = 5$ or $N \geq 7$. The seemingly obvious property ULNC for $\mathrm{PG}(N) \subset \mathbf{R}^2$ is difficult to prove. We opt for a less direct approach, instead verifying that no weak tangent of $\mathrm{PG}(N)$ contains any nontrivial geodesic segments. The result then follows from Theorems 5.7 and 4.5.

We begin by identifying the weak tangents of $\mathrm{PG}(N)$. Denote by $\mathscr{P}$ the family of all regular N-gons in $\mathbf{R}^2$. For each P in $\mathscr{P}$, let ϕ_P be a similarity map from P onto $P_\emptyset$. The regular N-gons $\phi_P^{-1}(P_j)$, $j = 1, \ldots, N$, are called the *children* of P. A regular N-gon P' is called a *descendant* of another regular N-gon P if there is a sequence $P = P(0), P(1), \ldots, P(k) = P'$ of regular N-gons so that $P(j)$ is a child of $P(j-1)$ for each $j = 1, \ldots, k$, or if $P' = P$.

The following is a variation on a definition of David and Semmes [5, Chapter 13].

Definition 6.1. A family $\mathscr{C}$ of regular N-gons is a *normalized family* if

(i) $\mathscr{C}$ contains the initial N-gon $P_\emptyset$,
(ii) each element of $\mathscr{C}$ is the child of exactly one element of $\mathscr{C}$,
(iii) all children of a given element of $\mathscr{C}$ are in $\mathscr{C}$,
(iv) each pair of elements of $\mathscr{C}$ are descended from a common ancestor.

For each $m \in \mathbf{Z}$, denote by $\widehat{K}(\mathscr{C})_m$ the union of the N-gons in $\mathscr{C}$ with side length r_N^m. The *limiting set* $\widehat{K}(\mathscr{C})$ for the family $\mathscr{C}$ is

$$\widehat{K}(\mathscr{C}) = \bigcap_{m \in \mathbf{Z}} \widehat{K}(\mathscr{C})_m.$$

David and Semmes [5, Definition 13.1] define a *normalized family of cubes* to be a family of closed cubes in $\mathbf{R}^n$ satisfying conditions (i), (ii), (iv) and

(iii$'$) for each element Q of $\mathscr{C}$, at least one child of Q lies in $\mathscr{C}$.

Our condition (iii) severely limits the number and variety of normalized families. For example, when $N = 4$ there are only three normalized families of squares in $\mathbf{R}^2$, modulo isometry of the limiting set. In fact, for any normalized family of squares in $\mathbf{R}^2$, the limiting set is isometric with one of the following: the quadrant $\{(x,y) : x \ge 0,\ y \ge 0\}$, the half-plane $\{(x,y) : y \ge 0\}$, or the full plane $\mathbf{R}^2$.

Proposition 6.2. *Let $N \ge 3$ be fixed. Then, up to rescaling, each weak tangent of* $\mathrm{PG}(N)$ *is isometric with either* $\mathrm{PG}(N)$ *or with* $\widehat{\mathrm{PG}(N)}(\mathscr{C})$ *for some normalized family $\mathscr{C}$ of N-gons.*

The proof of Lemma 13.9 in [5] can be adapted to the current setting; we omit the details.

Modulo Proposition 6.2, the proof of Theorem 1.6 reduces to verifying that no limiting set $\widehat{\mathrm{PG}(N)}(\mathscr{C})$ contains a nontrivial line segment when $N = 5$ or $N \ge 7$. It suffices to prove that $\mathrm{PG}(N)$ itself contains no nontrivial line segments. To see this, observe that if $\widehat{\mathrm{PG}(N)}(\mathscr{C})$ is contained a nontrivial (compact) line segment L, then L is contained within one of the polygons P in $\mathscr{C}$, and hence can be mapped by the similarity ϕ_P to a nontrivial line segment in $\mathrm{PG}(N) \subset P_\emptyset$.

Proposition 6.3. *If $N = 5$ or $N \ge 7$, then* $\mathrm{PG}(N)$ *contains no nontrivial line segments.*

Proof. For $m = 0, 1, \ldots$ and $N \ge 3$, let $l_m(N)$ be the length of the longest line segment contained in K_m. Recall that K_m denotes the union of the level m polygons, i.e., mth generation children of $P_\emptyset$. Observe that $l_m(3) = 1$ and $l_m(4) = \sqrt{2}$ for all $m \ge 0$, while $l_0(6) = 2$ and $l_m(6) = \sqrt{3}$ for all $m \ge 1$. For each $N = 5, 7, 8, 9, \ldots$, we will show that

$$l_m(N) \le 12 r_N^m \operatorname{diam}(P_\emptyset) \tag{6.4}$$

for all m. Since $r_N < 1$, this suffices to complete the proof.

The post-critically finite case $N \not\equiv 0 \pmod 4$ is easier and will be treated first. In this case we have $P_i \cap P_{i+1} = \{v_i\}$ for $i = 0, \ldots, N-1$ (indices taken modulo N), and the junction points $v_0, \ldots, v_{N-1}$ form the vertices of a new regular N-gon. No three of these vertices are collinear, which implies that any line segment contained within K_1 meets at most three of the children of $P_\emptyset$. Thus

$$l_1(N) \le 3\,\mathrm{diam}(P_i) = 3r_N\,\mathrm{diam}(P_\emptyset).$$

In a similar manner, we prove that any line segment L contained within $K_2 \subset K_1$ meets at most six grandchildren of $P_\emptyset$. Since $N \ne 3, 6$, any line segment joining distinct level one junction vertices v_i, v_j is not entirely contained within K_2. Consequently L meets at most one such vertex, and is contained in two adjacent children of $P_\emptyset$. In this case, the grandchildren in question are necessarily children of at most two adjacent children of $P_\emptyset$. By the previous paragraph, L meets at most three children of each of these two children, whence

$$l_2(N) \le 6\,\mathrm{diam}(P_{ij}) = 6r_N^2\,\mathrm{diam}(P_\emptyset).$$

Continuing in this fashion, we can prove that any line segment contained within K_m meets at most 12 level m descendants of $P_\emptyset$. Thus

$$l_m(N) \le 12\,\mathrm{diam}(P_{i_1\cdots i_m}) = 12r_N^m\,\mathrm{diam}(P_\emptyset) \quad \text{for all } m.$$

Next, we consider the case $N \equiv 0 \pmod 4$. In this case we have $P_i \cap P_{i+1} = E_i$ for $i = 0, \ldots, N-1$ (indices taken modulo N), where E_i is a junction edge common to P_i and P_{i+1}. A somewhat involved geometric argument (which we omit) shows that at most three junction edges are collinear (we say that sets $A_1, \ldots, A_k$ are *collinear* if there exist points $a_j \in A_j$, $j = 1, \ldots, k$, so that $a_1, \ldots, a_k$ are collinear). Thus any line segment contained within K_1 meets at most four children of $P_\emptyset$, and

$$l_1(N) \le 4\,\mathrm{diam}(P_i) = 4r_N\,\mathrm{diam}(P_\emptyset).$$

Next, we prove that any line segment L contained within K_2 meets at most eight of the grandchildren of $P_\emptyset$. Since $N \ne 4$, any line segment joining two distinct level one junction edges E_i, E_j is not entirely contained within K_2. (This is another geometric argument which we omit.) Consequently, L meets at most one such edge, the grandchildren in question are necessarily children of at most two fixed adjacent children of $P_\emptyset$, and the desired conclusion follows from the previous paragraph. We conclude that

$$l_2(N) \le 8\,\mathrm{diam}(P_{ij}) = 8r_N^2\,\mathrm{diam}(P_\emptyset)$$

and similarly, that

$$l_m(N) \le 8\,\mathrm{diam}(P_{i_1\cdots i_m}) = 8r_N^m\,\mathrm{diam}(P_\emptyset)$$

for all m. With these computations the proof of the proposition is complete. □

7. Appendix

Upon reading an early draft of this paper, Tomi Laakso [8] suggested a further geometric characterization of snowflake spaces. The results and proofs in this section are due to Laakso.

Definition 7.1 (Laakso). A metric space (X,d) is *line-fitting* if for each $n\in\mathbf{N}$ there is a metric d_n on the disjoint union $X\coprod[0,1]$ such that (i) d_n is the usual Euclidean metric in $[0,1]$, (ii) d_n is a constant multiple of d in X and (iii) $[0,1]$ is contained in the $1/n$-neighborhood of X.

Informally, (X,d) is line-fitting if it is possible to "attach" line segments to X with arbitrarily fine accuracy.

The main result of this section is the following theorem.

Theorem 7.2. *(X,d) is a snowflake if and only if it is not line-fitting.*

An advantage of this characterization is that it does not require any auxiliary embeddings of the space in "nice" Banach spaces. However, while the concept of line-fitting is an intrinsic geometric notion, it may be rather complicated to verify that a given space is not line fitting. Furthermore, the estimates for the exponent $p>1$ which arise in the following proof are less precise than those obtained in the setting of Theorem 1.5, when the metric space in question admits a suitable embedding. See Theorem 3.1.

Sketch of proof. It is easy to prove that a snowflake space cannot be line-fitting. Indeed, suppose that (X,d) is a p-snowflake for some $p>1$ and is line-fitting. Choose $c>0$ so that (1.2) holds, choose metrics d_n, $n\in\mathbf{N}$, on $X\coprod[0,1]$, and associated scaling factors r_n so that $d_n|_X=r_nd$ as in Definition 7.1, and consider the points $x_i=i/n$, $i=0,\dots,n$. By the hypothesis, there are points y_i, $i=0,\dots,n$ in X so that $d_n(x_i,y_i)<1/n$ for all i, whence

$$cr_n^p\left(1-\frac{2}{n}\right)^p<cr_n^pd(y_0,y_n)^p\le\sum_{i=1}^n r_n^pd(y_i,y_{i-1})^p<r_n^pn\left(\frac{3}{n}\right)^p.$$

For sufficiently large n, this gives a contradiction.

Conversely, suppose that X is not line-fitting, and choose $s>0$ so that for any metric d' on $X\coprod[0,1]$ with d' a constant multiple of d on X and d' equal to the Euclidean metric on $[0,1]$, the estimate

$$\sup_{0\le t\le 1}\operatorname{dist}(t,X)>s \tag{7.3}$$

holds. Choose J and $\delta>0$ so that $2^{-J}<2/s$ and $(1+2\delta)^J<1+s$.

For any x and y in X, consider the following recursive procedure. Begin with $p_0=x$ and $p_1=y$. Assume that points $p_{i\cdot 2^{-j}}$, $i=0,\dots,2^j$, are given for some $j\ge 0$. If one of the lens sets $L_i:=L\big(p_{(i-1)2^{-j}},p_{i\cdot 2^{-j}};\frac12,\delta\big)$, $i=1,\dots,2^j$, is empty, terminate the procedure. Otherwise, for each such i, choose $p_{(2i-1)2^{-j-1}}\in L_i$.

Claim. *This procedure terminates in fewer than J steps.*

To see why the claim holds, suppose that this procedure may be continued at least J steps, yielding points $p_{i\cdot 2^{-J}}$, $i=0,\dots,2^J$, in X. From the construction, it is easy to see that the estimate

$$2^{-J}|i-l| - s \le \frac{d(p_{i\cdot 2^{-J}}, p_{l\cdot 2^{-J}})}{d(x,y)} \le 2^{-J}|i-l| + s \tag{7.4}$$

holds true for each $0 \le i, l \le 2^J$. Set $r = d(x,y)^{-1}$ and define a distance function d' in $X \coprod [0,1]$ as follows: $d'|_X = rd$, $d'|_{[0,1]} =$ Euclidean metric, and

$$d'(a,t) = \min_{i=0,\dots,2^J} rd(a, p_{i\cdot 2^{-J}}) + |i\cdot 2^{-J} - t| + \tfrac{1}{2}s$$

for all $a \in X$ and $t \in [0,1]$. Observe that $d'(p_{i\cdot 2^{-J}}, i\cdot 2^{-J}) = \frac{1}{2}s$ for each $i = 0,\dots,2^J$. From (7.4) it easily follows that d' is a metric on $X \coprod [0,1]$ and $[0,1]$ is contained in the $\frac{3}{4}s$-neighborhood of X. This contradicts (7.3).

If one could conclude from the claim that

$$L(x,y;\lambda,\xi) = \emptyset \tag{7.5}$$

for some $0 < \lambda < 1$ and $\xi > 0$, then the snowflake property would follow as in the proof of Theorem 3.1. However, (7.5) need not be true in general. Nevertheless, (X,d) is a p-snowflake as may be shown by an elaboration on the proof of Theorem 3.1.

It suffices to prove that (1.2) holds for some $c > 0$ and $p > 1$. Let $\beta = \{x_0, x_1, \dots, x_N\}$ be a chain of points in X with $N \ge 2^J$. (If $N \le 2^J$ the conclusion follows with any p and $c = 2^{-pJ}$ by the triangle inequality.) If $\beta \cap L\big(x_0, x_N; \frac{1}{2}, \delta\big) = \emptyset$ then (1.2) holds for a suitable choice of $p > 1$ and $c > 0$ depending only on δ as in the proof of Theorem 3.1. Otherwise, choose $x_k \in \beta \cap L\big(x_0, x_N; \frac{1}{2}, \delta\big)$ and consider the subchains $\beta_L = \{x_0,\dots,x_k\}$ and $\beta_R = \{x_k,\dots,x_N\}$. If $\beta_L \cap L\big(x_0, x_k; \frac{1}{2}, \delta\big) = \emptyset$ then $\sum_{i=1}^k d(x_i, x_{i-1})^p \ge cd(x_0,x_k)^p$, while if $\beta_R \cap L\big(x_k, x_N; \frac{1}{2}, \delta\big) = \emptyset$, then $\sum_{i=k+1}^N d(x_i,x_{i-1})^p \ge cd(x_k,x_N)^p$ for the same choice of $p > 1$ and $c > 0$ as above. If both sets are nonempty, then choose $x_m \in \beta_L \cap L\big(x_0, x_k; \frac{1}{2}, \delta\big)$ and $x_q \in \beta_R \cap L\big(x_k, x_N; \frac{1}{2}, \delta\big)$ and iterate the above procedure.

By the claim, this procedure will terminate in fewer than J steps, and yield indices $i', i'' \in \{0,\dots,N\}$, $i' \le i''$, so that $\tilde\beta \cap L\big(x_{i'}, x_{i''}; \frac{1}{2}, \delta\big) = \emptyset$, where $\tilde\beta = \{x_{i'},\dots,x_{i''}\}$. Again it follows from the proof of Theorem 3.1 that

$$\sum_{i=i'+1}^{i''} d(x_i,x_{i-1})^p \ge cd(x_{i'}, x_{i''})^p.$$

The estimate $d(x_{i'}, x_{i''}) \geq \left(\frac{1}{2}-\delta\right)^J d(x_0, x_N)$ follows from the construction, whence

$$\sum_{i=1}^{N} d(x_i, x_{i-1})^p \geq \sum_{i=i'+1}^{i''} d(x_i, x_{i-1})^p \geq c\left(\tfrac{1}{2}-\delta\right)^J d(x_0, x_N)^p.$$

Thus (1.2) holds for some $p > 1$ and $c' = c\left(\frac{1}{2}-\delta\right)^J > 0$ depending only on s. (Recall that J and δ were chosen depending only on s.) This completes the proof. □

References

[1] Assouad, P.: Plongements lipschitziens dans $\mathbf{R}^n$. - Bull. Soc. Math. France 111, 1983, 429–448.

[2] Bonk, M., J. Heinonen, and P. Koskela: Uniformizing Gromov hyperbolic spaces. - Astérisque 270, 2001.

[3] Bonk, M., and O. Schramm: Embeddings of Gromov hyperbolic spaces. - Geom. Funct. Anal. 10 (2), 2000, 266–306.

[4] Burago, D., Y. Burago, and S. Ivanov: A Course in Metric Geometry. - Grad. Stud. Math. 33, American Mathematical Society, Providence, RI, 2001.

[5] David, G., and S. Semmes: Fractured Fractals and Broken Dreams: Self-Similar Geometry through Metric and Measure. - Oxford Lecture Ser. Math. Appl. 7, Oxford University Press, 1997.

[6] Heinonen, J.: Lectures on Analysis on Metric Spaces. - Springer-Verlag, New York, 2001.

[7] Keith, S., and Laakso, T.: Conformal Assouad dimension and modulus. - Geom. Funct. Anal. (to appear).

[8] Laakso, T.: Personal communication, 2004.

[9] Laakso, T. J.: Plane with A_∞-weighted metric not bi-Lipschitz embeddable to $\mathbf{R}^N$. - Bull. London Math. Soc. 34, 2002, 667–676.

[10] Lemin, A. J.: Isometric embedding of ultrametric (non-Archimedean) spaces in Hilbert space and Lebesgue space. - In: p-Adic Functional Analysis (Ioannina, 2000), Lecture Notes in Pure and Appl. Math. 222, Dekker, New York, 2001, 203–218.

[11] Semmes, S.: On the nonexistence of bi-Lipschitz parameterizations and geometric problems about A_∞-weights. - Rev. Mat. Iberoamericana 12, 1996, 337–410.

[12] Tyson, J. T., and J. M. Wu: Quasiconformal dimensions of self-similar fractals. - Preprint, 2003. Available at `http://www.math.uiuc.edu/~tyson/papers/qcdim.pdf`.

Received 10 June 2004
Received in revised form 8 October 2004