An invitation to Alexandrov geometry: CAT(0) spaces

Stephanie Alexander,
Vitali Kapovitch,
Anton Petrunin

Contents

Preface

This short monograph arose as an offshoot of the book on Alexandrov geometry [9] we have been writing for a number of years. The notes were shaped in a number of lectures given by the third author to undergraduate students at different occasions in Penn State, including the MASS program, at the Summer School “Algebra and Geometry” in Yaroslavl, and at SPbSU.

The idea is to demonstrate the beauty and power of Alexandrov geometry by reaching interesting applications and theorems with a minimum of preparation. Namely, we consider CAT(0) spaces — the metric spaces with nonpositive curvature in the sense of Alexandrov; these spaces can be loosely described as a non-linear generalization of a Hilbert space.

In Lecture 1, we discuss necessary preliminaries.

In Lecture 2, we discuss the Reshetnyak gluing theorem and apply it to a problem in billiards which was solved by Dmitri Burago, Serge Ferleger, and Alexey Kononenko.

In Lecture 3 we apply Lecture 2 to a problem in billiards which was solved by Dmitri Burago, Serge Ferleger, and Alexey Kononenko.

Lecture 4 provides the so-called Reshetnyak's majorization theorem. It is illustrated by several applications about convexity and geodesics.

In Lecture 5, we discuss the Hadamard–Cartan globalization theorem,.

In Lecture 6 we apply Lecture 5 to the construction of exotic aspherical manifolds introduced by Michael Davis.

In Lecture 7, we discuss examples of Alexandrov spaces with curvature bounded above. It is is based largely on work of Samuel Shefel on nonsmooth saddle surfaces.

Finally, in Lecture 8 we discuss barycenters with their applications to the dimension theory.

Here is a list of some sources providing a good introduction to

Alexandrov spaces with curvature bounded above, which we recommend for further information; we will not assume familiarity with any of these sources.

- ◇ The book by Martin Bridson and André Haefliger [25];
- ◇ Lecture notes of Werner Ballmann [19];
- ◇ Chapter 9 in the book [27] by Dmitri Burago, Yuri Burago and Sergei Ivanov.
- ◇ Our book [9].

Early history of Alexandov geometry

The idea that the essence of curvature lies in a condition on quadruples of points apparently originated with Abraham Wald. It is found in his publication on "coordinate-free differential geometry" [90] written under the supervision of Karl Menger; the story of this discovery can be found in [69]. In 1941, similar definitions were rediscovered independently by Alexandr Danilovich Alexandrov, see [12]. In Alexandrov's work the first fruitful applications of this approach were given. Mainly:

- ◇ Alexandrov's embedding theorem — *metrics of non-negative curvature on the sphere, and only they, are isometric to closed convex surfaces in Euclidean 3-space.*
- ◇ Gluing theorem, which tells when the sphere obtained by gluing of two discs along their boundaries has non-negative curvature in the sense of Alexandrov.

These two results together gave a very intuitive geometric tool for studying embeddings and bending of surfaces in Euclidean space, and changed this subject dramatically. They formed the foundation of the branch of geometry now called Alexandrov geometry.

The study of spaces with curvature bounded above started later. The first paper on the subject was written by Alexandrov; it appeared in 1951, see [14]. It was based on work of Herbert Busemann, who studied spaces satisfying a weaker condition [32].

Yurii Grigorievich Reshetnyak proved fundamental results about general spaces with curvature bounded above, the most important of which is his gluing and majorization theorems. An equally important theorem is the Hadamard–Cartan theorem (globalization theorem). These theorems and their history are discussed in lectures 2, 4 and 5.

Surfaces with upper curvature bounds were studied extensively in the 50-s and 60-s, and are by now well understood; see the survey [80] and the references therein.

Manifesto of Alexandrov geometry

Alexandrov geometry can use "back to Euclid" as a slogan. Alexandrov spaces are defined via axioms similar to those given by Euclid, but certain equalities are changed to inequalities. Depending on the sign of the inequalities, we get Alexandrov spaces with curvature bounded above or curvature bounded below. The definitions of the two classes of spaces are similar, but their properties and known applications are quite different.

Consider the space $\mathcal{M}_4$ of all isometry classes of 4-point metric spaces. Each element in $\mathcal{M}_4$ can be described by 6 numbers — the distances between all 6 pairs of its points, say $\ell_{i,j}$ for $1 \leqslant i < j \leqslant 4$ modulo permutations of the index set $(1, 2, 3, 4)$. These 6 numbers are subject to 12 triangle inequalities; that is,

$$\ell_{i,j} + \ell_{j,k} \geqslant \ell_{i,k}$$

holds for all i, j and k, where we assume that $\ell_{j,i} = \ell_{i,j}$ and $\ell_{i,i} = 0$.

Consider the subset $\mathcal{E}_4 \subset \mathcal{M}_4$ of all isometry classes of 4-point metric spaces that admit isometric embeddings into Euclidean space. The complement $\mathcal{M}_4 \setminus \mathcal{E}_4$ has two connected components.

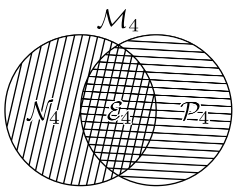

0.1. Exercise. *Prove the latter statement.*

One of the components will be denoted by $\mathcal{P}_4$ and the other by $\mathcal{N}_4$. Here $\mathcal{P}$ and $\mathcal{N}$ stand for positive and negative curvature because spheres have no quadruples of type $\mathcal{N}_4$ and hyperbolic space has no quadruples of type $\mathcal{P}_4$.

A metric space, with length metric, that has no quadruples of points of type $\mathcal{P}_4$ or $\mathcal{N}_4$ respectively is called an Alexandrov space with non-positive or non-negative curvature.

Here is an exercise, solving which would force the reader to rebuild a considerable part of Alexandrov geometry. It might be helpful to spend some time thinking about this exercise before proceeding. (The length metric is defined in Section 1C.)

0.2. Advanced exercise. *Assume $\mathcal{X}$ is a complete metric space with length metric, containing only quadruples of type $\mathcal{E}_4$. Show that $\mathcal{X}$ is isometric to a convex set in a Hilbert space.*

In the definition above, instead of Euclidean space one can take hyperbolic space of curvature -1. In this case, one obtains the definition of spaces with curvature bounded above or below by -1.

To define spaces with curvature bounded above or below by 1, one has to take the unit 3-sphere and specify that only the quadruples of points such that each of the four triangles has perimeter less than $2\cdot\pi$ are checked. The latter condition could be considered as a part of the spherical triangle inequality.

Acknowledgment

We want to thank David Berg, Richard Bishop, Yurii Burago, Maxime Fortier Bourque, Sergei Ivanov, Michael Kapovich, Bernd Kirchheim, Bruce Kleiner, Nikolai Kosovsky, Greg Kuperberg, Nina Lebedeva, John Lott, Alexander Lytchak, Dmitri Panov, Stephan Stadler, Wilderich Tuschmann, and Sergio Zamora Barrera for a number of discussions and suggestions.

We thank the mathematical institutions where we worked on this material, including BIRS, MFO, Henri Poincaré Institute, University of Colone, Max Planck Institute for Mathematics.

The first author was partially supported by the Simons Foundation grant #209053. The second author was partially supported by a Discovery grant from NSERC and by the Simons Foundation grant #390117. The third author was partially supported by the NSF grants DMS 1309340 and DMS-2005279 and the Simons Foundation #584781.

Lecture 1

Preliminaries

In this lecture we fix some conventions and recall the main definitions. It may be used as a quick reference when reading the book.

To learn background in metric geometry, the reader may consult the book of Dmitri Burago, Yuri Burago, and Sergei Ivanov [27] or the book by the third author [74].

A Metric spaces

The distance between two points x and y in a metric space $\mathcal{X}$ will be denoted by $|x-y|$ or $|x-y|_{\mathcal{X}}$. The latter notation is used if we need to emphasize that the distance is taken in the space $\mathcal{X}$.

The function

$$\operatorname{dist}_x\colon y\mapsto |x-y|$$

is called the *distance function* from x.

- The *diameter* of a metric space $\mathcal{X}$ is defined as

$$\operatorname{diam}\mathcal{X}=\sup\{\,|x-y|_{\mathcal{X}}\ :\ x,y\in\mathcal{X}\,\}.$$

- Given $R\in[0,\infty]$ and $x\in\mathcal{X}$, the sets

$$\begin{aligned}\mathrm{B}(x,R)&=\{y\in\mathcal{X}\mid |x-y|<R\},\\ \overline{\mathrm{B}}[x,R]&=\{y\in\mathcal{X}\mid |x-y|\leqslant R\}\end{aligned}$$

are called, respectively, the *open* and the *closed balls* of radius R with center x. Again, if we need to emphasize that these balls are taken in the metric space $\mathcal{X}$, we write

$$\mathrm{B}(x,R)_{\mathcal{X}}\quad\text{and}\quad\overline{\mathrm{B}}[x,R]_{\mathcal{X}}.$$

B Geodesics, triangles, and hinges

Geodesic. Let $\mathcal{X}$ be a metric space and $\mathbb{I}$ be a real interval. A globally isometric map $\gamma\colon \mathbb{I}\to\mathcal{X}$ is called a geodesic[1]; in other words, $\gamma\colon\mathbb{I}\to\mathcal{X}$ is a geodesic if

$$|\gamma(s)-\gamma(t)|_{\mathcal{X}} = |s-t|$$

for any pair $s,t\in\mathbb{I}$.

We say that $\gamma\colon\mathbb{I}\to\mathcal{X}$ is a geodesic from point p to point q if $\mathbb{I}=[a,b]$ and $p=\gamma(a)$, $q=\gamma(b)$. In this case the image of γ is denoted by $[pq]$ and with an abuse of notations we also call it a geodesic.

We may write $[pq]_{\mathcal{X}}$ to emphasize that the geodesic $[pq]$ is in the space $\mathcal{X}$. We also use the following shortcut notation:

$$]pq[\,=[pq]\setminus\{p,q\},\qquad]pq]=[pq]\setminus\{p\},\qquad [pq[\,=[pq]\setminus\{q\}.$$

In general, a geodesic between p and q need not exist and if it exists, it need not be unique. However, once we write $[pq]$ we mean that we have made a choice of geodesic.

A metric space is called geodesic if any pair of its points can be joined by a geodesic.

A geodesic path is a geodesic with constant-speed parametrization by $[0,1]$.

A curve $\gamma\colon\mathbb{I}\to\mathcal{X}$ is called a local geodesic if for any $t\in\mathbb{I}$ there is a neighborhood U of t in $\mathbb{I}$ such that the restriction $\gamma|_U$ is a geodesic. A constant-speed parametrization of a local geodesic by the unit interval $[0,1]$ is called a local geodesic path.

Triangle. For a triple of points $p,q,r\in\mathcal{X}$, a choice of a triple of geodesics $([qr],[rp],[pq])$ will be called a triangle; we will use the short notation $[pqr]=([qr],[rp],[pq])$.

Again, given a triple $p,q,r\in\mathcal{X}$ there may be no triangle $[pqr]$ simply because one of the pairs of these points cannot be joined by a geodesic. Also, many different triangles with these vertices may exist, any of which can be denoted by $[pqr]$. However, if we write $[pqr]$, it means that we have made a choice of such a triangle; that is, we have fixed a choice of the geodesics $[qr]$, $[rp]$, and $[pq]$.

The value

$$|p-q|+|q-r|+|r-p|$$

will be called the perimeter of the triangle $[pqr]$.

Hinge. Let $p,x,y\in\mathcal{X}$ be a triple of points such that p is distinct from x and y. A pair of geodesics $([px],[py])$ will be called a hinge and will be denoted by $[p\,^x_y]$.

[1]Various authors call it differently: shortest path, minimizing geodesic.

Convex set. A set A in a metric space $\mathcal{X}$ is called *convex* if for every two points $p, q \in A$, *every* geodesic $[pq]$ in $\mathcal{X}$ lies in A.

A set $A \subset \mathcal{X}$ is called *locally convex* if every point $a \in A$ admits an open neighborhood $\Omega \ni a$ in $\mathcal{X}$ such that any geodesic lying in Ω and with ends in A lies completely in A.

Note that any open set is locally convex by definition.

C Length spaces

A *curve* is defined as a continuous map from a real interval to a space. If the real interval is $[0, 1]$, then the curve is called a *path*.

1.1. Definition. *Let $\mathcal{X}$ be a metric space and $\alpha\colon \mathbb{I} \to \mathcal{X}$ be a curve. We define the length of α as*

$$\operatorname{length}\alpha := \sup_{t_0 \leqslant t_1 \leqslant \ldots \leqslant t_n} \sum_i |\alpha(t_i) - \alpha(t_{i-1})|.$$

Directly from the definition, it follows that if a path $\alpha\colon [0,1] \to \mathcal{X}$ connects two points x and y (that is, if $\alpha(0) = x$ and $\alpha(1) = y$), then

$$\operatorname{length}\alpha \geqslant |x - y|.$$

Let A be a subset of a metric space $\mathcal{X}$. Given two points $x, y \in A$, consider the value

$$|x - y|_A = \inf_\alpha \{\operatorname{length}\alpha\},$$

where the infimum is taken for all paths α from x to y in A.[2]

If $|x - y|_A$ takes finite value for each pair $x, y \in A$, then $|x - y|_A$ defines a metric on A; this metric will be called the *induced length metric* on A.

If for any $\varepsilon > 0$ and any pair of points x and y in a metric space $\mathcal{X}$, there is a path α connecting x to y such that

$$\operatorname{length}\alpha < |x - y| + \varepsilon,$$

then $\mathcal{X}$ is called a *length space* and the metric on $\mathcal{X}$ is called a *length metric*.

If $f\colon \tilde{\mathcal{X}} \to \mathcal{X}$ is a covering, then a length metric on $\mathcal{X}$ can be lifted to $\tilde{\mathcal{X}}$ by declaring

$$\operatorname{length}_{\tilde{\mathcal{X}}} \gamma = \operatorname{length}_{\mathcal{X}}(f \circ \gamma)$$

[2]Note that while this notation slightly conflicts with the previously defined notation for distance on a general metric space, we will usually work with ambient length spaces where the meaning will be unambiguous.

for any curve γ in $\tilde{\mathcal{X}}$. The space $\tilde{\mathcal{X}}$ with this metric is called the metric cover of $\mathcal{X}$.

Note that any geodesic space is a length space. As can be seen from the following example, the converse does not hold.

1.2. Example. *Let $\mathcal{X}$ be obtained by gluing a countable collection of disjoint intervals $\{\mathbb{I}_n\}$ of length $1+\frac{1}{n}$, where for each $\mathbb{I}_n$ the left end is glued to p and the right end to q. Then $\mathcal{X}$ carries a natural complete length metric with respect to which $|p-q|=1$ but there is no geodesic connecting p to q.*

1.3. Exercise. *Give an example of a complete length space for which no pair of distinct points can be joined by a geodesic.*

Let $\mathcal{X}$ be a metric space and $x, y \in \mathcal{X}$.

(i) A point $z \in \mathcal{X}$ is called a midpoint between x and y if

$$|x-z| = |y-z| = \tfrac{1}{2}\cdot|x-y|.$$

(ii) Assume $\varepsilon \geqslant 0$. A point $z \in \mathcal{X}$ is called an ε-midpoint between x and y if

$$|x-z|, \quad |y-z| \leqslant \tfrac{1}{2}\cdot|x-y| + \varepsilon.$$

Note that a 0-midpoint is the same as a midpoint.

1.4. Menger's lemma. *Let $\mathcal{X}$ be a complete metric space.*

(a) Assume that for any pair of points $x, y \in \mathcal{X}$ and any $\varepsilon > 0$ there is an ε-midpoint z. Then $\mathcal{X}$ is a length space.

(b) Assume that for any pair of points $x, y \in \mathcal{X}$, there is a midpoint z. Then $\mathcal{X}$ is a geodesic space.

The second part of this lemma was proved by Karl Menger [69, Section 6].

Proof. We first prove *(a)*. Let $x, y \in \mathcal{X}$ be a pair of points.

Set $\varepsilon_n = \frac{\varepsilon}{4^n}$, $\alpha(0) = x$ and $\alpha(1) = y$.

Let $\alpha(\frac{1}{2})$ be an ε_1-midpoint between $\alpha(0)$ and $\alpha(1)$. Further, let $\alpha(\frac{1}{4})$ and $\alpha(\frac{3}{4})$ be ε_2-midpoints between the pairs $(\alpha(0), \alpha(\frac{1}{2}))$ and $(\alpha(\frac{1}{2}), \alpha(1))$ respectively. Applying the above procedure recursively, on the n-th step we define $\alpha(\frac{k}{2^n})$, for every odd integer k such that $0 < \frac{k}{2^n} < 1$, as an ε_n-midpoint between the already defined $\alpha(\frac{k-1}{2^n})$ and $\alpha(\frac{k+1}{2^n})$.

In this way we define $\alpha(t)$ for $t \in W$, where W denotes the set of dyadic rationals in $[0,1]$. Since $\mathcal{X}$ is complete, the map α can be extended continuously to $[0,1]$. Moreover,

❶
$$\begin{aligned}\operatorname{length}\alpha &\leqslant |x-y| + \sum_{n=1}^{\infty} 2^{n-1}\cdot\varepsilon_n \leqslant \\ &\leqslant |x-y| + \tfrac{\varepsilon}{2}.\end{aligned}$$

Since $\varepsilon > 0$ is arbitrary, we get *(a)*.

To prove *(b)*, one should repeat the same argument taking midpoints instead of ε_n-midpoints. In this case ❶ holds for $\varepsilon_n = \varepsilon = 0$. □

A metric space $\mathcal{X}$ is called proper if all closed bounded sets in $\mathcal{X}$ are compact. This condition is equivalent to each of the following statements:

1. For some (and therefore any) point $p \in \mathcal{X}$ and any $R < \infty$, the closed ball $\overline{\mathrm{B}}[p,R] \subset \mathcal{X}$ is compact.
2. The function $\operatorname{dist}_p \colon \mathcal{X} \to \mathbb{R}$ is proper for some (and therefore any) point $p \in \mathcal{X}$; that is, for any compact set $K \subset \mathbb{R}$, its inverse image $\{\, x \in \mathcal{X} : |p-x|_{\mathcal{X}} \in K \,\}$ is compact.

Since in a compact space a sequence of $\frac{1}{n}$-midpoints z_n contains a convergent subsequence, Menger's lemma immediately implies the following.

1.5. Proposition. *A proper length space is geodesic.*

1.6. Hopf–Rinow theorem. *Any complete, locally compact length space is proper.*

Proof. Let $\mathcal{X}$ be a locally compact length space. Given $x \in \mathcal{X}$, denote by $\rho(x)$ the supremum of all $R > 0$ such that the closed ball $\overline{\mathrm{B}}[x,R]$ is compact. Since $\mathcal{X}$ is locally compact,

❷
$$\rho(x) > 0 \quad \text{for any} \quad x \in \mathcal{X}.$$

It is sufficient to show that $\rho(x) = \infty$ for some (and therefore any) point $x \in \mathcal{X}$.

Assume the contrary; that is, $\rho(x) < \infty$. We claim that

❸ $B = \overline{\mathrm{B}}[x, \rho(x)]$ *is compact for any* x.

Indeed, $\mathcal{X}$ is a length space; therefore for any $\varepsilon > 0$, the set $\overline{\mathrm{B}}[x, \rho(x) - \varepsilon]$ is a compact ε-net in B. Since B is closed and hence complete, it must be compact. △

Next we claim that

❹ $|\rho(x) - \rho(y)| \leqslant |x-y|_{\mathcal{X}}$ *for any* $x, y \in \mathcal{X}$*; in particular* $\rho\colon \mathcal{X} \to \mathbb{R}$ *is a continuous function.*

Indeed, assume the contrary; that is, $\rho(x) + |x - y| < \rho(y)$ for some $x, y \in \mathcal{X}$. Then $\overline{\mathrm{B}}[x, \rho(x) + \varepsilon]$ is a closed subset of $\overline{\mathrm{B}}[y, \rho(y)]$ for some $\varepsilon > 0$. Then compactness of $\overline{\mathrm{B}}[y, \rho(y)]$ implies compactness of $\overline{\mathrm{B}}[x, \rho(x) + \varepsilon]$, a contradiction. △

Set $\varepsilon = \min\{\,\rho(y) \,:\, y \in B\,\}$; the minimum is defined since B is compact. From ❷, we have $\varepsilon > 0$.

Choose a finite $\frac{\varepsilon}{10}$-net $\{a_1, a_2, \dots, a_n\}$ in B. The union W of the closed balls $\overline{\mathrm{B}}[a_i, \varepsilon]$ is compact. Clearly $\overline{\mathrm{B}}[x, \rho(x) + \frac{\varepsilon}{10}] \subset W$. Therefore $\overline{\mathrm{B}}[x, \rho(x) + \frac{\varepsilon}{10}]$ is compact, a contradiction. □

1.7. Exercise. *Construct a geodesic space that is locally compact, but whose completion is neither geodesic nor locally compact.*

D Constructions

Product space. Given two metric spaces $\mathcal{U}$ and $\mathcal{V}$, the product space $\mathcal{U} \times \mathcal{V}$ is defined as the set of all pairs (u, v) where $u \in \mathcal{U}$ and $v \in \mathcal{V}$ with the metric defined by formula

$$|(u^1, v^1) - (u^2, v^2)|_{\mathcal{U}\times\mathcal{V}} = \sqrt{|u^1 - u^2|^2_{\mathcal{U}} + |v^1 - v^2|^2_{\mathcal{V}}}.$$

1.8. Exercise. *Show that product of length spaces is a length space.*

1.9. Exercise. *Show that projection of a geodesic path from the product space to its factors are geodesic paths.*

Cone. The cone $\mathcal{V} = \mathrm{Cone}\,\mathcal{U}$ over a metric space $\mathcal{U}$ is defined as the metric space whose underlying set consists of equivalence classes in $[0, \infty) \times \mathcal{U}$ with the equivalence relation "$\sim$" given by $(0, p) \sim (0, q)$ for any points $p, q \in \mathcal{U}$, and whose metric is given by the cosine rule

$$|(p, s) - (q, t)|_{\mathcal{V}} = \sqrt{s^2 + t^2 - 2{\cdot}s{\cdot}t{\cdot}\cos\alpha},$$

where $\alpha = \min\{\pi, |p - q|_{\mathcal{U}}\}$.

The point in the cone $\mathcal{V}$ formed by the equivalence class of $0 \times \mathcal{U}$ is called the tip of the cone and is denoted by 0 or $0_{\mathcal{V}}$. The distance $|0 - v|_{\mathcal{V}}$ is called the norm of v and is denoted by $|v|$ or $|v|_{\mathcal{V}}$.

1.10. Exercise. *Let $[pq]$ be a geodesic in* $\mathrm{Cone}\,\mathcal{U}$*; assume it does not pass thru the tip. Show that the projection of $[pq]$ to $\mathcal{U}$ (after reparametrization) is a geodesic of length less than π.*

Suspension. The suspension $\mathcal{V} = \operatorname{Susp}\mathcal{U}$ over a metric space $\mathcal{U}$ is defined as the metric space whose underlying set consists of equivalence classes in $[0,\pi]\times\mathcal{U}$ with the equivalence relation "$\sim$" given by $(0,p)\sim$ $\sim(0,q)$ and $(\pi,p)\sim(\pi,q)$ for any points $p,q\in\mathcal{U}$, and whose metric is given by the spherical cosine rule

$$\cos|(p,s)-(q,t)|_{\operatorname{Susp}\mathcal{U}} = \cos s\cdot\cos t - \sin s\cdot\sin t\cdot\cos\alpha,$$

where $\alpha = \min\{\pi, |p-q|_{\mathcal{U}}\}$.

The points in $\mathcal{V}$ formed by the equivalence classes of $0\times\mathcal{U}$ and $\pi\times\mathcal{U}$ are called the north and the south poles of the suspension.

1.11. Exercise. *Let $\mathcal{U}$ be a metric space. Show that the spaces*

$$\mathbb{R}\times\operatorname{Cone}\mathcal{U} \quad \textit{and} \quad \operatorname{Cone}[\operatorname{Susp}\mathcal{U}]$$

are isometric.

E Model angles and triangles

Let $\mathcal{X}$ be a metric space and $p,q,r\in\mathcal{X}$. Let us define the model triangle $[\tilde p\tilde q\tilde r]$ (briefly, $[\tilde p\tilde q\tilde r] = \tilde\triangle(pqr)_{\mathbb{E}^2}$) to be a triangle in the plane $\mathbb{E}^2$ with the same side lengths; that is,

$$|\tilde p-\tilde q| = |p-q|, \quad |\tilde q-\tilde r| = |q-r|, \quad |\tilde r-\tilde p| = |r-p|.$$

In the same way we can define the hyperbolic and the spherical model triangles $\tilde\triangle(pqr)_{\mathbb{H}^2}$, $\tilde\triangle(pqr)_{\mathbb{S}^2}$ in the hyperbolic plane $\mathbb{H}^2$ and the unit sphere $\mathbb{S}^2$. In the latter case the model triangle is said to be defined if in addition

$$|p-q|+|q-r|+|r-p| < 2\cdot\pi.$$

In this case the model triangle again exists and is unique up to an isometry of $\mathbb{S}^2$.

If $[\tilde p\tilde q\tilde r] = \tilde\triangle(pqr)_{\mathbb{E}^2}$ and $|p-q|,|p-r|>0$, the angle measure of $[\tilde p\tilde q\tilde r]$ at $\tilde p$ will be called the model angle of the triple p, q, r and will be denoted by $\tilde\measuredangle(p\,{}^q_r)_{\mathbb{E}^2}$. In the same way we define $\tilde\measuredangle(p\,{}^q_r)_{\mathbb{H}^2}$ and $\tilde\measuredangle(p\,{}^q_r)_{\mathbb{S}^2}$; in the latter case we assume in addition that the model triangle $\tilde\triangle(pqr)_{\mathbb{S}^2}$ is defined.

We may use the notation $\tilde\measuredangle(p\,{}^q_r)$ if it is evident which of the model spaces $\mathbb{H}^2$, $\mathbb{E}^2$ or $\mathbb{S}^2$ is meant.

1.12. Alexandrov's lemma. *Let p, x, y, z be distinct points in a metric space such that $z \in \,]xy[$. Then the following expressions for the Euclidean model angles have the same sign:*

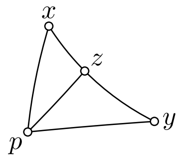

(a) $\tilde\measuredangle(x\,{}^{p}_{y}) - \tilde\measuredangle(x\,{}^{p}_{z})$,

(b) $\tilde\measuredangle(z\,{}^{p}_{x}) + \tilde\measuredangle(z\,{}^{p}_{y}) - \pi$.

Moreover,

$$\tilde\measuredangle(p\,{}^{x}_{y}) \geqslant \tilde\measuredangle(p\,{}^{x}_{z}) + \tilde\measuredangle(p\,{}^{z}_{y}),$$

with equality if and only if the expressions in (a) and (b) vanish.

The same holds for the hyperbolic and spherical model angles, but in the latter case one has to assume in addition that

$$|p-z| + |p-y| + |x-y| < 2\cdot\pi.$$

Proof. Consider the model triangle $[\tilde x\tilde p\tilde z] = \tilde\triangle(xpz)$. Take a point $\tilde y$ on the extension of $[\tilde x\tilde z]$ beyond $\tilde z$ so that $|\tilde x - \tilde y| = |x-y|$ (and therefore $|\tilde x - \tilde z| = |x - z|$).

Since increasing the opposite side in a plane triangle increases the corresponding angle, the following expressions have the same sign:

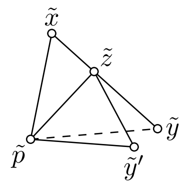

(i) $\measuredangle[\tilde x\,{}^{\tilde p}_{\tilde y}] - \tilde\measuredangle(x\,{}^{p}_{y})$,

(ii) $|\tilde p - \tilde y| - |p - y|$,

(iii) $\measuredangle[\tilde z\,{}^{\tilde p}_{\tilde y}] - \tilde\measuredangle(z\,{}^{p}_{y})$.

Since

$$\measuredangle[\tilde x\,{}^{\tilde p}_{\tilde y}] = \measuredangle[\tilde x\,{}^{\tilde p}_{\tilde z}] = \tilde\measuredangle(x\,{}^{p}_{z})$$

and

$$\measuredangle[\tilde z\,{}^{\tilde p}_{\tilde y}] = \pi - \measuredangle[\tilde z\,{}^{\tilde x}_{\tilde p}] = \pi - \tilde\measuredangle(z\,{}^{x}_{p}),$$

the first statement follows.

For the second statement, construct a model triangle $[\tilde p\tilde z\tilde y'] = \tilde\triangle(pzy)_{\mathbb{E}^2}$ on the opposite side of $[\tilde p\tilde z]$ from $[\tilde x\tilde p\tilde z]$. Note that

$$\begin{aligned}|\tilde x - \tilde y'| &\leqslant |\tilde x - \tilde z| + |\tilde z - \tilde y'| = \\ &= |x - z| + |z - y| = \\ &= |x - y|.\end{aligned}$$

Therefore

$$\begin{aligned}\tilde\measuredangle(p\,{}^{x}_{z}) + \tilde\measuredangle(p\,{}^{z}_{y}) &= \measuredangle[\tilde p\,{}^{\tilde x}_{\tilde z}] + \measuredangle[\tilde p\,{}^{\tilde z}_{\tilde y'}] = \\ &= \measuredangle[\tilde p\,{}^{\tilde x}_{\tilde y'}] \leqslant \\ &\leqslant \tilde\measuredangle(p\,{}^{x}_{y}).\end{aligned}$$

Equality holds if and only if $|\tilde x - \tilde y'| = |x - y|$, as required. □

F Angles and the first variation

Given a hinge $[p\,{}^x_y]$, we define its *angle* as the limit

$$❶ \qquad \measuredangle[p\,{}^x_y] := \lim_{\bar x,\bar y\to p} \tilde\measuredangle(p\,{}^{\bar x}_{\bar y})_{\mathbb{E}^2},$$

where $\bar x \in\,]px]$ and $\bar y \in\,]py]$. (The angle $\measuredangle[p\,{}^x_y]$ is defined if the limit exists.)

The value under the limit can be calculated from the cosine law:

$$\cos\tilde\measuredangle(p\,{}^x_y)_{\mathbb{E}^2} = \frac{|p-x|^2+|p-y|^2-|x-y|^2}{2\cdot|p-x|\cdot|p-y|}.$$

The following lemma implies that in ❶, one can use $\tilde\measuredangle(p\,{}^{\bar x}_{\bar y})_{\mathbb{S}^2}$ or $\tilde\measuredangle(p\,{}^{\bar x}_{\bar y})_{\mathbb{H}^2}$ instead of $\tilde\measuredangle(p\,{}^{\bar x}_{\bar y})_{\mathbb{E}^2}$.

1.13. Lemma. *For any three points p, x, y in a metric space the following inequalities*

$$❷ \qquad \begin{aligned} |\tilde\measuredangle(p\,{}^x_y)_{\mathbb{S}^2} - \tilde\measuredangle(p\,{}^x_y)_{\mathbb{E}^2}| &\leqslant |p-x|\cdot|p-y|,\\ |\tilde\measuredangle(p\,{}^x_y)_{\mathbb{H}^2} - \tilde\measuredangle(p\,{}^x_y)_{\mathbb{E}^2}| &\leqslant |p-x|\cdot|p-y| \end{aligned}$$

hold whenever the left-hand side is defined.

Proof. Note that

$$\tilde\measuredangle(p\,{}^x_y)_{\mathbb{H}^2} \leqslant \tilde\measuredangle(p\,{}^x_y)_{\mathbb{E}^2} \leqslant \tilde\measuredangle(p\,{}^x_y)_{\mathbb{S}^2}.$$

Therefore

$$\begin{aligned} 0 &\leqslant \tilde\measuredangle(p\,{}^x_y)_{\mathbb{S}^2} - \tilde\measuredangle(p\,{}^x_y)_{\mathbb{H}^2} \leqslant \\ &\leqslant \quad \tilde\measuredangle(p\,{}^x_y)_{\mathbb{S}^2} + \tilde\measuredangle(x\,{}^p_y)_{\mathbb{S}^2} + \tilde\measuredangle(y\,{}^p_x)_{\mathbb{S}^2} - \\ &\quad - \tilde\measuredangle(p\,{}^x_y)_{\mathbb{H}^2} - \tilde\measuredangle(x\,{}^p_y)_{\mathbb{H}^2} - \tilde\measuredangle(y\,{}^p_x)_{\mathbb{H}^2} = \\ &= \operatorname{area}\tilde\triangle(pxy)_{\mathbb{S}^2} + \operatorname{area}\tilde\triangle(pxy)_{\mathbb{H}^2}. \end{aligned}$$

The inequality ❷ follows since

$$\begin{aligned} 0 &\leqslant \operatorname{area}\tilde\triangle(pxy)_{\mathbb{H}^2} \leqslant \\ &\leqslant \operatorname{area}\tilde\triangle(pxy)_{\mathbb{S}^2} \leqslant \\ &\leqslant |p-x|\cdot|p-y|. \end{aligned}$$

□

1.14. Triangle inequality for angles. *Let $[px^1]$, $[px^2]$ and $[px^3]$ be three geodesics in a metric space. If all the angles $\alpha^{ij} = \measuredangle[p\,{}^{x^i}_{x^j}]$ are defined, then they satisfy the triangle inequality:*

$$\alpha^{13} \leqslant \alpha^{12} + \alpha^{23}.$$

Proof. Since $\alpha^{13} \leqslant \pi$, we may assume that $\alpha^{12} + \alpha^{23} < \pi$.

Let γ^i be the unit-speed parametrization of $[px^i]$ from p to x^i. Given any $\varepsilon > 0$, for all sufficiently small $t, \tau, s \in \mathbb{R}_+$ we have

$$\begin{aligned}|\gamma^1(t) - \gamma^3(\tau)| \leqslant |\gamma^1(t) - \gamma^2(s)| + |\gamma^2(s) - \gamma^3(\tau)| <\\ < \sqrt{t^2 + s^2 - 2{\cdot}t{\cdot}s{\cdot}\cos(\alpha^{12} + \varepsilon)} +\\ + \sqrt{s^2 + \tau^2 - 2{\cdot}s{\cdot}\tau{\cdot}\cos(\alpha^{23} + \varepsilon)} \leqslant\end{aligned}$$

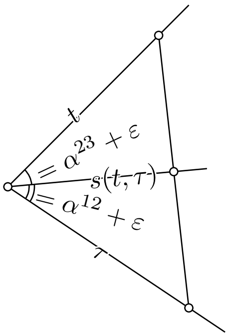

Below we define $s(t, \tau)$ so that for $s = s(t, \tau)$, this chain of inequalities can be continued as follows:

$$\leqslant \sqrt{t^2 + \tau^2 - 2{\cdot}t{\cdot}\tau{\cdot}\cos(\alpha^{12} + \alpha^{23} + 2{\cdot}\varepsilon)}.$$

Thus for any $\varepsilon > 0$,

$$\alpha^{13} \leqslant \alpha^{12} + \alpha^{23} + 2{\cdot}\varepsilon.$$

Hence the result.

To define $s(t, \tau)$, consider three rays $\tilde\gamma^1$, $\tilde\gamma^2$, $\tilde\gamma^3$ on a Euclidean plane starting at one point, such that $\measuredangle(\tilde\gamma^1, \tilde\gamma^2) = \alpha^{12} + \varepsilon$, $\measuredangle(\tilde\gamma^2, \tilde\gamma^3) = \alpha^{23} + \varepsilon$ and $\measuredangle(\tilde\gamma^1, \tilde\gamma^3) = \alpha^{12} + \alpha^{23} + 2{\cdot}\varepsilon$. We parametrize each ray by the distance from the starting point. Given two positive numbers $t, \tau \in \mathbb{R}_+$, let $s = s(t, \tau)$ be the number such that $\tilde\gamma^2(s) \in [\tilde\gamma^1(t)\ \tilde\gamma^3(\tau)]$. Clearly $s \leqslant \max\{t, \tau\}$, so t, τ, s may be taken sufficiently small. $\square$

1.15. Exercise. *Prove that the sum of adjacent angles is at least π.*

More precisely, let $\mathcal{X}$ be a complete length space and $p, x, y, z \in \mathcal{X}$. If $p \in\ ]xy[$, then

$$\measuredangle[p\,{}^x_z] + \measuredangle[p\,{}^y_z] \geqslant \pi$$

whenever each angle on the left-hand side is defined.

1.16. First variation inequality. *Assume that for a hinge $[q\,{}^p_x]$ the angle $\alpha = \measuredangle[q\,{}^p_x]$ is defined. Then*

$$|p - \gamma(t)| \leqslant |q - p| - t{\cdot}\cos\alpha + o(t),$$

where γ is the unit-speed parametrization of $[qx]$ from q to x.

Proof. Take a sufficiently small $\varepsilon > 0$. Denote by β the the unit-speed parametrization of $[qp]$ from q to p. For all sufficiently small $t > 0$, we have

$$\begin{aligned}|\beta(t/\varepsilon) - \gamma(t)| \leqslant \tfrac{t}{\varepsilon}{\cdot}\sqrt{1 + \varepsilon^2 - 2{\cdot}\varepsilon{\cdot}\cos\alpha} + o(t) \leqslant\\ \leqslant \tfrac{t}{\varepsilon} - t{\cdot}\cos\alpha + t{\cdot}\varepsilon.\end{aligned}$$

Applying the triangle inequality, we get

$$\begin{aligned}|p-\gamma(t)| &\leqslant |p-\beta(t/\varepsilon)|+|\beta(t/\varepsilon)-\gamma(t)| \leqslant \\ &\leqslant |p-q|-t\cdot\cos\alpha+t\cdot\varepsilon\end{aligned}$$

for any fixed $\varepsilon>0$ and all sufficiently small t. Hence the result. $\square$

G Space of directions and tangent space

Let $\mathcal{X}$ be a metric space with defined angles for all hinges. Fix a point $p\in\mathcal{X}$.

Consider the set $\mathfrak{S}_p$ of all nontrivial geodesics that start at p. By 1.14, the triangle inequality holds for the angle measure $\measuredangle$ on $\mathfrak{S}_p$, so $(\mathfrak{S}_p,\measuredangle)$ forms a pseudometric space; that is, $\measuredangle$ satisfies all the conditions of a metric on $\mathfrak{S}_p$, except that the angle between distinct geodesics might vanish.

The metric space corresponding to $(\mathfrak{S}_p,\measuredangle)$ is called the space of geodesic directions at p, denoted by Σ'_p or $\Sigma'_p\mathcal{X}$. Elements of Σ'_p are called geodesic directions at p. Each geodesic direction is formed by an equivalence class of geodesics in $\mathfrak{S}_p$ for the equivalence relation

$$[px]\sim[py] \quad\Longleftrightarrow\quad \measuredangle[p\,{}^x_y]=0.$$

The completion of Σ'_p is called the space of directions at p and is denoted by Σ_p or $\Sigma_p\mathcal{X}$. Elements of Σ_p are called directions at p.

The Euclidean cone $\operatorname{Cone}\Sigma_p$ over the space of directions Σ_p is called the tangent space at p and is denoted by T_p or $\mathrm{T}_p\mathcal{X}$.

The tangent space T_p could also be defined directly, without introducing the space of directions. To do so, consider the set $\mathfrak{T}_p$ of all geodesics with constant-speed parametrizations starting at p. Given $\alpha,\beta\in\mathfrak{T}_p$, set

$$\text{❶}\qquad |\alpha-\beta|_{\mathfrak{T}_p}=\lim_{\varepsilon\to0}\frac{|\alpha(\varepsilon)-\beta(\varepsilon)|_{\mathcal{X}}}{\varepsilon}$$

Since the angles in $\mathcal{X}$ are defined, ❶ defines a pseudometric on $\mathfrak{T}_p$.

The corresponding metric space admits a natural isometric identification with the cone $\mathrm{T}'_p=\operatorname{Cone}\Sigma'_p$. The elements of T'_p are equivalence classes for the relation

$$\alpha\sim\beta \quad\Longleftrightarrow\quad |\alpha(t)-\beta(t)|_{\mathcal{X}}=o(t).$$

The completion of T'_p is therefore naturally isometric to T_p.

Elements of T_p will be called tangent vectors at p, regardless of the fact that T_p is only a metric cone and need not be a vector space. Elements of T'_p will be called geodesic tangent vectors at p.

H Hausdorff convergence

Let $\mathcal{X}$ be a metric space and $A \subset \mathcal{X}$. We will denote by $\operatorname{dist}_A(x)$ the distance from A to a point x in $\mathcal{X}$; that is,

$$\operatorname{dist}_A(x) := \inf \{ |a - x|_{\mathcal{X}} \ : \ a \in A \}.$$

It is natural to assume that $\operatorname{dist}_{\varnothing}(x) = \infty$ for any x.

1.17. Definition of Hausdorff convergence. *Given a sequence of closed sets $(A_n)_{n=1}^{\infty}$ in a metric space $\mathcal{X}$, a closed set $A_\infty \subset \mathcal{X}$ is called the Hausdorff limit of $(A_n)_{n=1}^{\infty}$, briefly $A_n \to A_\infty$, if*

$$\operatorname{dist}_{A_n}(x) \to \operatorname{dist}_{A_\infty}(x) \quad as \quad n \to \infty$$

for every $x \in \mathcal{X}$.

In this case, the sequence of closed sets $(A_n)_{n=1}^{\infty}$ is said to converge in the sense of Hausdorff.

Examples. Let D_n be the disc in the coordinate plane with center $(0, n)$ and radius n. Then D_n converges to the upper half-plane as $n \to \infty$.

Note that sequence of one-point sets $\{(0, n)\}$ converges to the empty set. Indeed, for any $\operatorname{dist}_{\{(0,n)\}}(x) \to \infty = \operatorname{dist}_{\varnothing}(x)$ for any x.

1.18. Exercise. *Let $A_n \to A_\infty$ as in Definition 1.17.*

Show that A_∞ is the set of all points p such that $p_n \to p$ for some sequence of points $p_n \in A_n$.

Does the converse hold? That is, suppose $(A_n)_{n=1}^{\infty}, A_\infty$ are closed sets such that A_∞ is the set of all points p such that $p_n \to p$ for some sequence of points $p_n \in A_n$. Does this imply that $A_n \to A_\infty$?

1.19. First selection theorem. *Let $\mathcal{X}$ be a proper metric space and $(A_n)_{n=1}^{\infty}$ be a sequence of closed sets in $\mathcal{X}$. Then the sequence $(A_n)_{n=1}^{\infty}$ has a convergent subsequence in the sense of Hausdorff.*

Proof. Since $\mathcal{X}$ is proper, there is a countable dense set $\{x_1, x_2, \dots\}$ in $\mathcal{X}$. We can assume that the sequence $d_n = \operatorname{dist}_{A_n}(x_k)$ is bounded for each k. Otherwise there is a subsequence such that $\operatorname{dist}_{A_n}(x) \to \infty$ as $n \to \infty$ for some (and therefore any) x. Therefore, the corresponding subsequence of A_n converges to the empty set.

Therefore, passing to a subsequence of $(A_n)_{n=1}^{\infty}$, we can assume that $\operatorname{dist}_{A_n}(x_k)$ converges as $n \to \infty$ for any fixed k.

Note that for each n, the function $\operatorname{dist}_{A_n} : \mathcal{X} \to \mathbb{R}$ is 1-Lipschitz and nonnegative. Therefore the sequence $\operatorname{dist}_{A_n}$ converges pointwise to a 1-Lipschitz nonnegative function $f \colon \mathcal{X} \to \mathbb{R}$.

Set $A_\infty = f^{-1}(0)$. Let us show that

$$\mathrm{dist}_{A_\infty}(y) \leqslant f(y)$$

for any y. Assume the contrary; that is, $f(z) < R < \mathrm{dist}_{A_\infty}(z)$ for some $z \in \mathcal{X}$ and $R > 0$. Then for any sufficiently large n, there is a point $z_n \in A_n$ such that $|x - z_n| \leqslant R$. Since $\mathcal{X}$ is proper, we can pass to a partial limit z_∞ of z_n as $n \to \infty$.

It is clear that $f(z_\infty) = 0$; that is, $z_\infty \in A_\infty$. (Note that this implies that $A_\infty \neq \varnothing$.) On the other hand,

$$\mathrm{dist}_{A_\infty}(y) \leqslant |z_\infty - y| \leqslant R < \mathrm{dist}_{A_\infty}(y),$$

a contradiction.

On the other hand, since f is 1-Lipschitz, $\mathrm{dist}_{A_\infty}(y) \geqslant f(y)$. Therefore

$$\mathrm{dist}_{A_\infty}(y) = f(y)$$

for any $y \in \mathcal{X}$. Hence the result. □

I Gromov–Hausdorff convergence

1.20. Definition. *Let $\{\, \mathcal{X}_\alpha : \alpha \in \mathcal{A} \,\}$ be a collection of metric spaces. A metric ρ on the disjoint union*

$$\boldsymbol{X} = \bigsqcup_{\alpha \in \mathcal{A}} \mathcal{X}_\alpha$$

is called a compatible metric if the restriction of ρ to every $\mathcal{X}_\alpha$ coincides with the original metric on $\mathcal{X}_\alpha$.

1.21. Definition. *Let $\mathcal{X}_1, \mathcal{X}_2, \dots$ and $\mathcal{X}_\infty$ be proper metric spaces and ρ be a compatible metric on their disjoint union $\boldsymbol{X}$. Assume that $\mathcal{X}_n$ is an open set in $(\boldsymbol{X}, \rho)$ for each $n \neq \infty$, and $\mathcal{X}_n \to \mathcal{X}_\infty$ in $(\boldsymbol{X}, \rho)$ as $n \to \infty$ in the sense of Hausdorff (see Definition 1.17).*

Then we say ρ defines a convergence[3] in the sense of Gromov–Hausdorff, and write $\mathcal{X}_n \to \mathcal{X}_\infty$ or $\mathcal{X}_n \xrightarrow{\rho} \mathcal{X}_\infty$. The space $\mathcal{X}_\infty$ is called the limit space of the sequence $(\mathcal{X}_n)$ along ρ.

Usually Gromov–Hausdorff convergence is defined differently. We prefer this definition since it induces convergence for a sequence of points $x_n \in \mathcal{X}_n$ (Exercise 1.18), as well as weak convergence of measures μ_n on $\mathcal{X}_n$, and so on, corresponding to convergence in the ambient space $(\boldsymbol{X}, \rho)$.

[3]Formally speaking, convergence in the topology induced by ρ on $\boldsymbol{X}$.

Once we write $\mathcal{X}_n \to \mathcal{X}_\infty$, we mean that we have made a choice of convergence. Note that for a fixed sequence of metric spaces $(\mathcal{X}_n)$, it might be possible to construct different Gromov–Hausdorff convergences, say $\mathcal{X}_n \xrightarrow{\rho} \mathcal{X}_\infty$ and $\mathcal{X}_n \xrightarrow{\rho'} \mathcal{X}'_\infty$, whose limit spaces $\mathcal{X}_\infty$ and $\mathcal{X}'_\infty$ need not be isometric to each other.

For example, for the constant sequence $\mathcal{X}_n \stackrel{iso}{=\!=} \mathbb{R}_{\geqslant 0}$ may converge to $\mathbb{R}_{\geqslant 0}$ and $\mathbb{R}$. The first convergence is evident and the second could be guessed from the diagram.

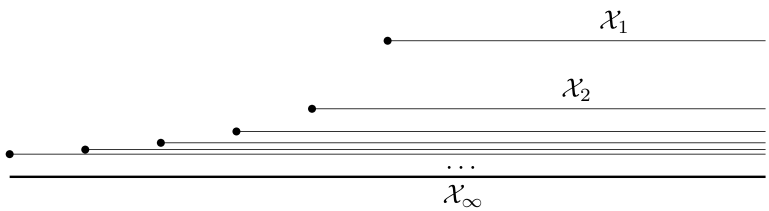

1.22. Second selection theorem. *Let $\mathcal{X}_n$ be a sequence of proper metric spaces with marked points $x_n \in \mathcal{X}_n$. Assume that for any fixed $R, \varepsilon > 0$, there is $N = N(R,\varepsilon) \in \mathbb{N}$ such that for each n the ball $\overline{\mathrm{B}}[x_n, R]_{\mathcal{X}_n}$ admits a finite ε-net with at most N points. Then there is a subsequence of $\mathcal{X}_n$ admitting a Gromov–Hausdorff convergence such that the sequence of marked points $x_n \in \mathcal{X}_n$ converges.*

Proof. From the main assumption in the theorem, in each space $\mathcal{X}_n$ there is a sequence of points $z_{i,n} \in \mathcal{X}_n$ such that the following condition holds for a fixed sequence of integers $M_1 < M_2 < \dots$

- ⋄ $|z_{i,n} - x_n|_{\mathcal{X}_n} \leqslant k+1$ if $i \leqslant M_k$,
- ⋄ the points $z_{1,n}, \dots, z_{M_k,n}$ form an $\tfrac{1}{k}$-net in $\overline{\mathrm{B}}[x_n, k]_{\mathcal{X}_n}$.

Passing to a subsequence, we can assume that the sequence

$$\ell_n = |z_{i,n} - z_{j,n}|_{\mathcal{X}_n}$$

converges for any i and j.

Consider a countable set of points $\mathcal{W} = \{w_1, w_2, \dots\}$ equipped with the pseudometric defined by

$$|w_i - w_j|_{\mathcal{W}} = \lim_{n\to\infty} |z_{i,n} - z_{j,n}|_{\mathcal{X}_n}.$$

Let $\hat{\mathcal{W}}$ be the metric space corresponding to $\mathcal{W}$; that is, points in $\hat{\mathcal{W}}$ are equivalence classes in $\mathcal{W}$ for the relation $\sim$, where $w_i \sim w_j$ if and only if $|w_i - w_j|_{\mathcal{W}} = 0$, and where

$$|[w_i] - [w_j]|_{\hat{\mathcal{W}}} := |w_i - w_j|_{\mathcal{W}}.$$

Denote by $\mathcal{X}_\infty$ the completion of $\hat{\mathcal{W}}$.

It remains to show that there is a Gromov–Hausdorff convergence $\mathcal{X}_n \to \mathcal{X}_\infty$ such that the sequence $x_n \in \mathcal{X}_n$ converges. To prove this, we need to construct a metric ρ on the disjoint union of

$$\boldsymbol{X} = \mathcal{X}_\infty \sqcup \mathcal{X}_1 \sqcup \mathcal{X}_2 \sqcup \dots$$

satisfying definitions 1.20 and 1.21. The metric ρ can be constructed as the maximal compatible metric such that

$$\rho(z_{i,n}, w_i) \leqslant \tfrac{1}{m}$$

for any $n \geqslant N_m$ and $i < I_m$ for a suitable choice of two sequences (I_m) and (N_m) with $I_1 = N_1 = 1$. $\square$

1.23. Exercise. *Let $\mathcal{K}$ be a compact metric space and*

$$f\colon \mathcal{K} \to \mathcal{K}$$

be a distance non-decreasing map. Prove that f is an isometry.

1.24. Exercise. *Let $\mathcal{X}_n$ be a sequence of metric spaces that admits two convergences $\mathcal{X}_n \xrightarrow{\rho} \mathcal{X}_\infty$ and $\mathcal{X}_n \xrightarrow{\rho'} \mathcal{X}'_\infty$.*

(a) If $\mathcal{X}_\infty$ is compact, then $\mathcal{X}_\infty \overset{iso}{=\!=} \mathcal{X}'_\infty$.

(b) If $\mathcal{X}_\infty$ is proper and there is a sequence of points $x_n \in \mathcal{X}_n$ that converges in both convergences, then $\mathcal{X}_\infty \overset{iso}{=\!=} \mathcal{X}'_\infty$.

J Remarks

It seems that Hausdorff convergence was first introduced by Felix Hausdorff in [54], and a couple of years later an equivalent definition was given by Wilhelm Blaschke in [22]. A refinement of this definition was introduced by Zdeněk Frolík in [46], and later rediscovered by Robert Wijsman in [91]. However, this refinement takes an intermediate place between the original Hausdorff convergence and closed convergence, also introduced by Hausdorff in [54]. For this reason we call it Hausdorff convergence instead of Hausdorff–Blascke–Frolík–Wijsman convergence.

Lecture 2

Gluing

In this lecture we define CAT(κ) spaces and prove Reshetnyak's gluing theorem.

Here "CAT" is an acronym for Cartan, Alexandrov, and Toponogov. It was coined by Mikhael Gromov in 1987. Originally, Alexandrov called these spaces "$\mathfrak{R}_\kappa$ domain"; this term is still in use.

A The 4-point condition

Given a quadruple of points p, q, x, y in a metric space $\mathcal{X}$, consider two model triangles in the plane $[\tilde{p}\tilde{x}\tilde{y}] = \tilde{\triangle}(pxy)_{\mathbb{E}^2}$ and $[\tilde{q}\tilde{x}\tilde{y}] = \tilde{\triangle}(qxy)_{\mathbb{E}^2}$ with common side $[\tilde{x}\tilde{y}]$.

If the inequality

$$|p-q|_{\mathcal{X}} \leqslant |\tilde{p}-\tilde{z}|_{\mathbb{E}^2} + |\tilde{z}-\tilde{q}|_{\mathbb{E}^2}$$

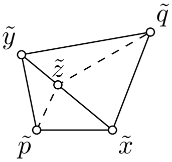

holds for any point $\tilde{z} \in [\tilde{x}\tilde{y}]$, then we say that the quadruple p, q, x, y satisfies CAT(0) comparison.

If we do the same for spherical model triangles $[\tilde{p}\tilde{x}\tilde{y}] = \tilde{\triangle}(pxy)_{\mathbb{S}^2}$ and $[\tilde{q}\tilde{x}\tilde{y}] = \tilde{\triangle}(qxy)_{\mathbb{S}^2}$, then we arrive at the definition of CAT(1) comparison. If one of the spherical model triangles is undefined,[1] then it is assumed that CAT(1) comparison automatically holds for this quadruple.

We can do the same for the model plane of curvature κ; that is, a sphere if $\kappa > 0$, Euclidean plane if $\kappa = 0$ and Lobachevsky plane if

[1] That is, if

$$|p-x| + |p-y| + |x-y| \geqslant 2\cdot\pi \quad \text{or} \quad |q-x| + |q-y| + |x-y| \geqslant 2\cdot\pi.$$

$\kappa < 0$. In this case we arrive at the definition of CAT(κ) comparison. However in these notes we will mostly consider CAT(0) comparison and occasionally CAT(1) comparison; so, if you see CAT(κ), you can assume that κ is 0 or 1.

If all quadruples in a metric space $\mathcal{X}$ satisfy CAT(κ) comparison, then we say that the space $\mathcal{X}$ is CAT(κ). (Note that CAT(κ) is an adjective.)

In order to check CAT(κ) comparison, it is sufficient to know the 6 distances between all pairs of points in the quadruple. This observation implies the following.

2.1. Proposition. *Any Gromov–Hausdorff limit of a sequence of* CAT(κ) *spaces is* CAT(κ).

In the proposition above, it does not matter which definition of convergence for metric spaces you use, as long as any quadruple of points in the limit space can be arbitrarily well approximated by quadruples in the sequence of metric spaces. In particular, it works for the so-called ultralimits.

B Geodesics

The CAT comparison can be applied to any metric space, but it is usually applied to geodesic spaces (or complete length spaces). To simplify the presentation we will often assume in addition that the space is proper. The latter means that any closed ball is compact.

Recall that function is proper if the inverse image of any compact set is compact. Note that *a metric space is proper if and only if the distance function from one (and therefore any) point is proper*; see Section 1A.

2.2. Proposition. *Let $\mathcal{U}$ be a complete geodesic* CAT(0) *space. Then any two points in $\mathcal{U}$ are joined by a unique geodesic.*

Proof. Suppose there are two geodesics between x and y. Then we can choose two points $p \neq q$ on these geodesics such that $|x-p| = |x-q|$ and therefore $|y-p| = |y-q|$.

Observe that the model triangles $[\tilde{p}\tilde{x}\tilde{y}] = \tilde{\triangle}(pxy)$ and $[\tilde{q}\tilde{x}\tilde{y}] = \tilde{\triangle}(qxy)$ are degenerate and moreover $\tilde{p} = \tilde{q}$. Applying CAT(0) comparison with $\tilde{z} = \tilde{p} = \tilde{q}$, we get that $|p-q| = 0$, a contradiction. $\square$

2.3. Exercise. *Given a hinge $[p\,{}^{x}_{y}]$ in a* CAT(0) *space $\mathcal{U}$, consider the function*

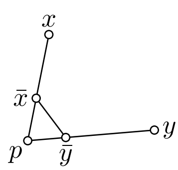

$$f\colon (|p-\bar{x}|, |p-\bar{y}|) \mapsto \tilde{\measuredangle}(p\,{}^{\bar{x}}_{\bar{y}}),$$

where $\bar{x} \in \,]px]$ *and* $\bar{y} \in \,]py]$. *Show that* f *is nondecreasing in each argument.*

Conclude that any hinge in a CAT(0) *space has defined angle.*

2.4. Exercise. *Fix a point* p *in a complete geodesic* CAT(0) *space* $\mathcal{U}$. *Given a point* $x \in \mathcal{U}$, *denote by* $\gamma_x\colon [0,1] \to \mathcal{U}$ *a (necessarily unique) geodesic path from* p *to* x.

Show that the family of maps $h_t\colon \mathcal{U} \to \mathcal{U}$ *defined by*

$$h_t(x) = \gamma_x(t)$$

is a homotopy. Conclude that $\mathcal{U}$ *is contractible.*

The homotopy in the exercise is a special case of the so-called geodesic homotopy. Namely, given two maps $h_0, h_1\colon \mathcal{X} \to \mathcal{U}$ we define homotopy

$$h_t(x) = \gamma_x(t),$$

where γ_x is the geodesic path from $h_0(x)$ to $h_1(x)$.

The construction in the previous exercise should help to solve the next one.

2.5. Exercise. *Let* $\mathcal{U}$ *be a complete geodesic* CAT(0) *space. Assume* $\mathcal{U}$ *is a topological manifold. Show that any geodesic in* $\mathcal{U}$ *can be extended as a two-side infinite geodesic.*

2.6. Exercise. *Assume* $\mathcal{U}$ *is a proper length* CAT(0) *space with extendable geodesics; that is, any geodesic is an arc in a local geodesic* $\mathbb{R} \to \mathcal{U}$. *Show that the space of geodesic directions at any point in* $\mathcal{U}$ *is complete.*

Does the statement remain true if $\mathcal{U}$ *is complete, but not required to be proper?*

C Thin triangles

Let $[\tilde{x}^1\tilde{x}^2\tilde{x}^3] = \tilde{\triangle}(x^1x^2x^3)$ be a model triangle for a triangle $[x^1x^2x^3]$ in a metric space. The map that sends a point $\tilde{z} \in [\tilde{x}^i\tilde{x}^j]$ to the corresponding point $z \in [x^ix^j]$ for each side will be called natural.

2.7. Definition. *A triangle* $[xyz]$ *in the metric space* $\mathcal{U}$ *is called thin if the natural map* $\tilde{\triangle}(xyz)_{\mathbb{E}^2} \to [xyz]$ *is distance nonincreasing.*

Analogously, a triangle $[xyz]$ *is called spherically thin if the natural map from the spherical model triangle* $\tilde{\triangle}(xyz)_{\mathbb{S}^2}$ *to* $[xyz]$ *is distance nonincreasing.*

2.8. Proposition. *A geodesic space is* CAT(0) *(*CAT(1)*) if and only if all its triangles are thin (respectively, all its triangles of perimeter* $< 2\cdot\pi$ *are spherically thin).*

Proof; if part. Apply the triangle inequality and thinness of triangles $[pxy]$ and $[qxy]$, where p, q, x, and y are as in the definition of the CAT(κ) comparison.

Only-if part. Applying CAT(0) comparison to a quadruple p, q, x, y with $q \in [xy]$ shows that any triangle satisfies point-side comparison, that is, the distance from a vertex to a point on the opposite side is no greater than the corresponding distance in the Euclidean model triangle.

Now consider a triangle $[xyz]$ and let $p \in [xy]$ and $q \in [xz]$. Let $\tilde p$, $\tilde q$ be the corresponding points on the sides of the model triangle $\tilde\triangle(xyz)_{\mathbb{E}^2}$. Applying 2.3, we get that

$$\tilde\measuredangle(x\,{}^y_z)_{\mathbb{E}^2} \geqslant \tilde\measuredangle(x\,{}^p_q)_{\mathbb{E}^2}.$$

Therefore $|\tilde p - \tilde q|_{\mathbb{E}^2} \geqslant |p-q|$.

The CAT(1) argument is the same. □

Recall that a curve $\gamma\colon \mathbb{I} \to \mathcal{U}$ is called a local geodesic if for any $t \in \mathbb{I}$ there is a neighborhood U of t in $\mathbb{I}$ such that the restriction $\gamma|_U$ is a geodesic.

2.9. Proposition. *Suppose $\mathcal{U}$ is a proper geodesic* CAT(0) *space. Then any local geodesic in $\mathcal{U}$ is a geodesic.*

Analogously, if $\mathcal{U}$ is a proper geodesic CAT(1) *space, then any local geodesic in $\mathcal{U}$ which is shorter than π is a geodesic.*

Proof. Suppose $\gamma\colon [0,\ell] \to \mathcal{U}$ is a local geodesic that is not a geodesic. Choose a to be the maximal value such that γ is a geodesic on $[0,a]$. Further, choose $b > a$ so that γ is a geodesic on $[a,b]$.

Since the triangle $[\gamma(0)\gamma(a)\gamma(b)]$ is thin (see the next section) and $|\gamma(0)-\gamma(b)| < b$ we have

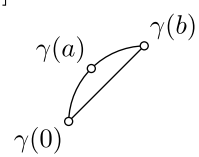

$$|\gamma(a-\varepsilon) - \gamma(a+\varepsilon)| < 2\cdot\varepsilon$$

for all small $\varepsilon > 0$. That is, γ is not length-minimizing on the interval $[a-\varepsilon, a+\varepsilon]$ for any $\varepsilon > 0$, a contradiction.

The spherical case is done in the same way. □

2.10. Exercise. *Let $\mathcal{U}$ be a complete geodesic space. Show that $\mathcal{U}$ is* CAT(0) *if and only if the function $f = \tfrac12\cdot\mathrm{dist}_p^2$ is 1-convex for* *any $p \in \mathcal{U}$; that is, the function $t \mapsto f \circ \gamma(t) - \frac{1}{2} \cdot t^2$ is convex for any unit-speed geodesic γ*

2.11. Exercise. *Suppose $\gamma_1, \gamma_2 \colon [0,1] \to \mathcal{U}$ are two geodesic paths in a complete geodesic* $\mathrm{CAT}(0)$ *space $\mathcal{U}$. Show that*

$$t \mapsto |\gamma_1(t) - \gamma_2(t)|_{\mathcal{U}}$$

is a convex function.

2.12. Exercise. *Let A be a convex closed set in a proper geodesic* $\mathrm{CAT}(0)$ *space $\mathcal{U}$; that is, if $x, y \in A$, then $[xy] \subset A$. Show that* dist_A *is convex.*

In particular, for any $r > 0$ the closed r-neighborhood of A is convex; that is, the set

$$A_r = \{\, x \in \mathcal{U} \;:\; \mathrm{dist}_A x \leqslant r \,\}$$

is convex.

2.13. Exercise. *Let $\mathcal{U}$ be a proper geodesic* $\mathrm{CAT}(0)$ *space and $K \subset \mathcal{U}$ be a closed convex set. Show that:*

(a) *For each point $p \in \mathcal{U}$ there is a unique point $p^* \in K$ that minimizes the distance $|p - p^*|$.*
(b) *The closest-point projection $p \mapsto p^*$ defined by (a) is short.*

Recall that a set A in a metric space $\mathcal{U}$ is called *locally convex* if for any point $p \in A$ there is an open neighborhood $\mathcal{U} \ni p$ such that any geodesic in $\mathcal{U}$ with ends in A lies in A.

2.14. Exercise. *Let $\mathcal{U}$ be a proper geodesic* $\mathrm{CAT}(0)$ *space. Show that any closed, connected, locally convex set in $\mathcal{U}$ is convex.*

D Inheritance lemma

2.15. Inheritance lemma. *Assume that a triangle $[pxy]$ in a metric space is decomposed into two triangles $[pxz]$ and $[pyz]$; that is, $[pxz]$ and $[pyz]$ have a common side $[pz]$, and the sides $[xz]$ and $[zy]$ together form the side $[xy]$ of $[pxy]$.*

If both triangles $[pxz]$ and $[pyz]$ are thin, then the triangle $[pxy]$ is also thin.

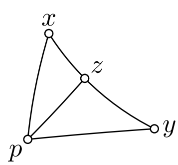

Analogously, if $[pxy]$ has perimeter $< 2 \cdot \pi$ and both triangles $[pxz]$ and $[pyz]$ are spherically thin, then triangle $[pxy]$ is spherically thin.

Proof. Construct the model triangles $[\dot p \dot x \dot z] = \tilde\triangle(pxz)_{\mathbb{E}^2}$ and $[\dot p \dot y \dot z] = \tilde\triangle(pyz)_{\mathbb{E}^2}$ so that $\dot x$ and $\dot y$ lie on opposite sides of $[\dot p \dot z]$.

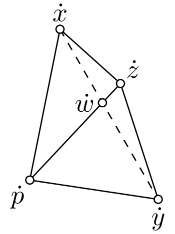

Let us show that

❶ $$\tilde\measuredangle(z\,{}^{p}_{x}) + \tilde\measuredangle(z\,{}^{p}_{y}) \geqslant \pi.$$

If not, then for some point $\dot w \in [\dot p\dot z]$, we have

$$|\dot x - \dot w| + |\dot w - \dot y| < |\dot x - \dot z| + |\dot z - \dot y| = |x - y|.$$

Let $w \in [pz]$ correspond to $\dot w$; that is, $|z - w| = |\dot z - \dot w|$. Since $[pxz]$ and $[pyz]$ are thin, we have

$$|x - w| + |w - y| < |x - y|,$$

contradicting the triangle inequality.

Denote by $\dot D$ the union of two solid triangles $[\dot p\dot x\dot z]$ and $[\dot p\dot y\dot z]$. Further, denote by $\tilde D$ the solid triangle $[\tilde p\tilde x\tilde y] = \tilde\triangle(pxy)_{\mathbb{E}^2}$. By ❶, there is a *short map*[2] $F\colon \tilde D \to \dot D$ that sends

$$\tilde p \mapsto \dot p, \qquad \tilde x \mapsto \dot x, \qquad \tilde z \mapsto \dot z, \qquad \tilde y \mapsto \dot y.$$

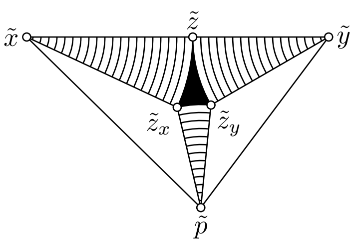

Indeed, by Alexandrov's lemma (1.12), there are nonoverlapping triangles

$$[\tilde p\tilde x\tilde z_x] \overset{iso}{=\!=} [\dot p\dot x\dot z]$$

and

$$[\tilde p\tilde y\tilde z_y] \overset{iso}{=\!=} [\dot p\dot y\dot z]$$

inside the triangle $[\tilde p\tilde x\tilde y]$.

Connect the points in each pair $(\tilde z, \tilde z_x)$, $(\tilde z_x, \tilde z_y)$ and $(\tilde z_y, \tilde z)$ with arcs of circles centered at $\tilde y$, $\tilde p$, and $\tilde x$ respectively. Define F as follows:

- ◇ Map $\mathrm{Conv}[\tilde p\tilde x\tilde z_x]$ isometrically onto $\mathrm{Conv}[\dot p\dot x\dot z]$; similarly map $\mathrm{Conv}[\tilde p\tilde y\tilde z_y]$ onto $\mathrm{Conv}[\dot p\dot y\dot z]$.
- ◇ If x is in one of the three circular sectors, say at distance r from its center, set $F(x)$ to be the point on the corresponding segment $[pz]$, $[xz]$ or $[yz]$ whose distance from the left-hand endpoint of the segment is r.
- ◇ Finally, if x lies in the remaining curvilinear triangle $\tilde z\tilde z_x\tilde z_y$, set $F(x) = z$.

By construction, F satisfies the conditions.

By assumption, the natural maps $[\dot p\dot x\dot z] \to [pxz]$ and $[\dot p\dot y\dot z] \to [pyz]$ are short. By composition, the natural map from $[\tilde p\tilde x\tilde y]$ to $[pyz]$ is short, as claimed.

The spherical case is done along the same lines. □

[2] In other words, distance-nonexpanding or 1-Lipschitz.

E Reshetnyak's gluing

Suppose $\mathcal{U}^1$ and $\mathcal{U}^2$ are proper geodesic spaces with isometric closed convex sets $A^i \subset \mathcal{U}^i$ and let $\iota\colon A^1 \to A^2$ be an isometry. Consider the space $\mathcal{W}$ of all equivalence classes in $\mathcal{U}^1 \sqcup \mathcal{U}^2$ with the equivalence relation given by $a \sim \iota(a)$ for any $a \in A^1$.

It is straightforward to see that $\mathcal{W}$ is a proper geodesic space when equipped with the following metric

$$\begin{aligned} |x-y|_{\mathcal{W}} &:= |x-y|_{\mathcal{U}^i} \\ &\qquad \text{if} \quad x,y \in \mathcal{U}^i, \quad \text{and} \\ |x-y|_{\mathcal{W}} &:= \min\left\{\, |x-a|_{\mathcal{U}^1} + |y-\iota(a)|_{\mathcal{U}^2} \;:\; a \in A^1 \,\right\} \\ &\qquad \text{if} \quad x \in \mathcal{U}^1 \quad \text{and} \quad y \in \mathcal{U}^2. \end{aligned}$$

Abusing notation, we denote by x and y the points in $\mathcal{U}^1 \sqcup \mathcal{U}^2$ and their equivalence classes in $\mathcal{U}^1 \sqcup \mathcal{U}^2/\!\sim$.

The space $\mathcal{W}$ is called the gluing of $\mathcal{U}^1$ and $\mathcal{U}^2$ along ι. If one applies this construction to two copies of one space $\mathcal{U}$ with a set $A \subset \mathcal{U}$ and the identity map $\iota\colon A \to A$, then the obtained space is called the double of $\mathcal{U}$ along A.

We can (and will) identify $\mathcal{U}^i$ with its image in $\mathcal{W}$; this way both subsets $A^i \subset \mathcal{U}^i$ will be identified and denoted further by A. Note that $A = \mathcal{U}^1 \cap \mathcal{U}^2 \subset \mathcal{W}$, therefore A is also a convex set in $\mathcal{W}$.

2.16. Reshetnyak gluing. *Suppose $\mathcal{U}^1$ and $\mathcal{U}^2$ are proper geodesic* CAT(0) *spaces with isometric closed convex sets $A^i \subset \mathcal{U}^i$, and $\iota\colon A^1 \to A^2$ is an isometry. Then the gluing of $\mathcal{U}^1$ and $\mathcal{U}^2$ along ι is a* CAT(0) *proper geodesic space.*

Proof. By construction of the gluing space, the statement can be reformulated in the following way:

2.17. Reformulation of 2.16. *Let $\mathcal{W}$ be a proper geodesic space with two closed convex sets $\mathcal{U}^1, \mathcal{U}^2 \subset \mathcal{W}$ such that $\mathcal{U}^1 \cup \mathcal{U}^2 = \mathcal{W}$ and $\mathcal{U}^1$, $\mathcal{U}^2$ are* CAT(0)*. Then $\mathcal{W}$ is* CAT(0).

It suffices to show that any triangle $[xyz]$ in $\mathcal{W}$ is thin. This is obviously true if all three points x, y, z lie in one of $\mathcal{U}^i$. Thus, without loss of generality, we may assume that $x \in \mathcal{U}^1$ and $y, z \in \mathcal{U}^2$.

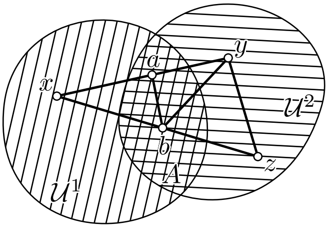

Choose points $a, b \in A = \mathcal{U}^1 \cap \mathcal{U}^2$ that lie respectively on the sides $[xy]$, $[xz]$. Note that

◇ the triangle $[xab]$ lies in $\mathcal{U}^1$,

◇ both triangles $[yab]$ and $[ybz]$ lie in $\mathcal{U}^2$.

In particular, each triangle $[xab]$, $[yab]$, and $[ybz]$ is thin.

Applying the inheritance lemma (2.15) twice, we get that $[xyb]$ and consequently $[xyz]$ is thin. □

Let A be a closed subset in a metric space $\mathcal{U}$. Gluing of two copies of $\mathcal{U}$ along the copies of A is called doubling $\mathcal{U}$ along A

2.18. Exercise. *Suppose $\mathcal{W}$ is a doubling of a geodesic space $\mathcal{U}$ along its closed subset A. Show that $\mathcal{W}$ is* CAT(0) *if and only if $\mathcal{U}$ is* CAT(0), *and A is convex in $\mathcal{U}$.*

F Comments

The gluing theorem (2.16) was proved by Yuri Reshetnyak [78]. It can be extended to all geodesic CAT(0) spaces. It also admits a natural generalization to geodesic CAT(κ) spaces; see the book of Martin Bridson and André Haefliger [25] and our book [9] for details.

Lecture 3

Billiards

A Puff pastry

In this section, we introduce the notion of Reshetnyak puff pastry. This construction will be used in the next section to prove the collision theorem (3.11).

Let $\boldsymbol{A} = (A^1, \dots, A^N)$ be an array of convex closed sets in the Euclidean space $\mathbb{E}^m$. Consider an array of $N+1$ copies of $\mathbb{E}^m$. Assume that the space $\mathcal{R}$ is obtained by gluing successive pairs of spaces along $A^1, \dots, A^N$ respectively.

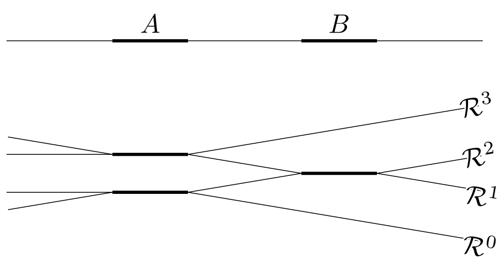

Puff pastry for (A, B, A).

The resulting space $\mathcal{R}$ will be called the Reshetnyak puff pastry for array $\boldsymbol{A}$. The copies of $\mathbb{E}^m$ in the puff pastry $\mathcal{R}$ will be called levels; they will be denoted by $\mathcal{R}^0, \dots, \mathcal{R}^N$. The point in the k-th level $\mathcal{R}^k$ that corresponds to $x \in \mathbb{E}^m$ will be denoted by x^k.

Given $x \in \mathbb{E}^m$, any point $x^k \in \mathcal{R}$ is called a lifting of x. The map $x \mapsto x^k$ defines an isometry $\mathbb{E}^m \to \mathcal{R}^k$; in particular, we can talk about liftings of subsets in $\mathbb{E}^m$.

Note that:

- ◇ The intersection $A^1 \cap \dots \cap A^N$ admits a unique lifting in $\mathcal{R}$.
- ◇ Moreover, $x^i = x^j$ for some $i < j$ if and only if

$$x \in A^{i+1} \cap \dots \cap A^j.$$

- ◇ For any k, the restriction $\mathcal{R}^k \to \mathbb{E}^m$ of the natural projection $x^k \mapsto x$ is an isometry.

3.1. Observation. *Any Reshetnyak puff pastry is a proper geodesic* CAT(0) *space.*

Proof. Apply Reshetnyak gluing theorem (2.16) recursively for the convex sets in the array. □

3.2. Proposition. *Assume $(A^1, \dots, A^N)$ and $(\check{A}^1, \dots, \check{A}^N)$ are two arrays of convex closed sets in $\mathbb{E}^m$ such that $A^k \subset \check{A}^k$ for each k. Let $\mathcal{R}$ and $\check{\mathcal{R}}$ be the corresponding Reshetnyak puff pastries. Then the map $\mathcal{R} \to \check{\mathcal{R}}$ defined by $x^k \mapsto \check{x}^k$ is short.*

Moreover, if

$$❶ \qquad |x^i - y^j|_{\mathcal{R}} = |\check{x}^i - \check{y}^j|_{\check{\mathcal{R}}}$$

for some $x, y \in \mathbb{E}^m$ and $i, j \in \{0, \dots, n\}$, then the unique geodesic $[\check{x}^i \check{y}^j]_{\check{\mathcal{R}}}$ is the image of the unique geodesic $[x^i y^j]_{\mathcal{R}}$ under the map $x^i \mapsto \check{x}^i$.

Proof. The first statement in the proposition follows from the construction of Reshetnyak puff pastries.

By Observation 3.1, $\mathcal{R}$ and $\check{\mathcal{R}}$ are proper geodesic CAT(0) spaces; hence $[x^i y^j]_{\mathcal{R}}$ and $[\check{x}^i \check{y}^j]_{\check{\mathcal{R}}}$ are unique. By ❶, since the map $\mathcal{R} \to \check{\mathcal{R}}$ is short, the image of $[x^i y^j]_{\mathcal{R}}$ is a geodesic of $\check{\mathcal{R}}$ joining $\check{x}^i$ to $\check{y}^j$. Hence the second statement follows. □

3.3. Definition. *Consider a Reshetnyak puff pastry $\mathcal{R}$ with the levels $\mathcal{R}^0, \dots, \mathcal{R}^N$. We say that $\mathcal{R}$ is end-to-end convex if $\mathcal{R}^0 \cup \mathcal{R}^N$, the union of its lower and upper levels, forms a convex set in $\mathcal{R}$; that is, if $x, y \in \mathcal{R}^0 \cup \mathcal{R}^N$, then $[xy]_{\mathcal{R}} \subset \mathcal{R}^0 \cup \mathcal{R}^N$.*

Note that if $\mathcal{R}$ is the Reshetnyak puff pastry for an array of convex sets $\boldsymbol{A} = (A^1, \dots, A^N)$, then $\mathcal{R}$ is end-to-end convex if and only if the union of the lower and the upper levels $\mathcal{R}^0 \cup \mathcal{R}^N$ is isometric to the double of $\mathbb{E}^m$ along the nonempty intersection $A^1 \cap \dots \cap A^N$.

3.4. Observation. *Let $\check{\boldsymbol{A}}$ and $\boldsymbol{A}$ be arrays of convex bodies in $\mathbb{E}^m$. Assume that array $\boldsymbol{A}$ is obtained by inserting in $\check{\boldsymbol{A}}$ several copies of the bodies which were already listed in $\check{\boldsymbol{A}}$.*

For example, if $\check{\boldsymbol{A}} = (A, C, B, C, A)$*, by placing* B *in the second place and* A *in the fourth place, we obtain* $\boldsymbol{A} = (A, B, C, A, B, C, A)$*.*

Denote by $\check{\mathcal{R}}$ *and* $\mathcal{R}$ *the Reshetnyak puff pastries for* $\check{\boldsymbol{A}}$ *and* $\boldsymbol{A}$ *respectively.*

If $\check{\mathcal{R}}$ *is end-to-end convex, then so is* $\mathcal{R}$*.*

Proof. Without loss of generality, we may assume that $\boldsymbol{A}$ is obtained by inserting one element in $\check{\boldsymbol{A}}$, say at the place number k.

Note that $\check{\mathcal{R}}$ is isometric to the puff pastry for $\boldsymbol{A}$ with A^k replaced by $\mathbb{E}^m$. It remains to apply Proposition 3.2. □

Let X be a convex set in a Euclidean space. By a dihedral angle, we understand an intersection of two half-spaces; the intersection of corresponding hyperplanes is called the edge of the angle. We say that a dihedral angle D supports X at a point $p \in X$ if D contains X and the edge of D contains p.

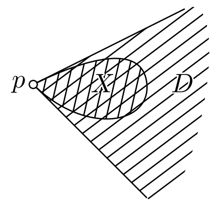

3.5. Lemma. *Let* A *and* B *be two-convex sets in* $\mathbb{E}^m$*. Assume that any dihedral angle supporting* $A \cap B$ *has angle measure at least* α*. Then the Reshetnyak puff pastry for the array*

$$(\underbrace{A, B, A, \dots}_{\lceil \frac{\pi}{\alpha} \rceil + 1 \textit{ times}}).$$

is end-to-end convex.

The proof of the lemma is based on a partial case, which we formulate as a sublemma.

3.6. Sublemma. *Let* $\ddot{A}$ *and* $\ddot{B}$ *be two half-planes in* $\mathbb{E}^2$*, where* $\ddot{A} \cap \ddot{B}$ *is an angle with measure* α*. Then the Reshetnyak puff pastry for the array*

$$(\underbrace{\ddot{A}, \ddot{B}, \ddot{A}, \dots}_{\lceil \frac{\pi}{\alpha} \rceil + 1 \textit{ times}})$$

is end-to-end convex.

Proof. Note that the puff pastry $\ddot{\mathcal{R}}$ is isometric to the cone over the space glued from the unit circles as shown on the diagram.

All the short arcs on the diagram have length α; the long arcs have length $\pi - \alpha$, so making a circuit along any path will take $2{\cdot}\pi$.

Applying Exercise 1.10, we get that the end-to-end convexity of $\ddot{\mathcal{R}}$ follows if any geodesic shorter than π with the ends on the inner and the outer circles lies completely in the union of these two circles. The latter holds if the zigzag line in the picture has length at least π. This line is formed by $\lceil\frac{\pi}{\alpha}\rceil$ arcs with length α each. Hence the sublemma. $\square$

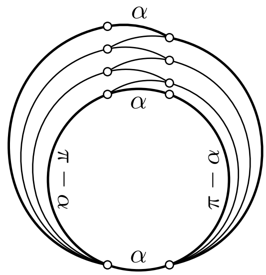

In the proof of 3.5, we will use the following exercise in convex geometry:

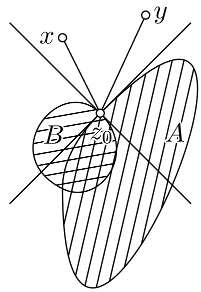

3.7. Exercise. *Let A and B be two closed convex sets in $\mathbb{E}^m$ and $A\cap B\neq\varnothing$. Given two points $x,y\in\mathbb{E}^m$ let $f(z)=|x-z|+|y-z|$.*

Let $z_0\in A\cap B$ be a point of minimum of $f|_{A\cap B}$.

Show that there are half-spaces $\dot{A}$ and $\dot{B}$ such that $\dot{A}\supset A$ and $\dot{B}\supset B$ and z_0 is also a point of minimum of the restriction $f|_{\dot{A}\cap\dot{B}}$.

Proof of 3.5. Fix arbitrary $x,y\in\mathbb{E}^m$. Choose a point $z\in A\cap B$ for which the sum

$$|x-z|+|y-z|$$

is minimal. To show the end-to-end convexity of $\mathcal{R}$, it is sufficient to prove the following:

❷ *The geodesic $[x^0y^N]_{\mathcal{R}}$ contains $z^0=z^N\in\mathcal{R}$.*

Without loss of generality, we may assume that $z\in\partial A\cap\partial B$. Indeed, since the puff pastry for the 1-array (B) is end-to-end convex, Proposition 3.2 together with 3.4 imply ❷ in case z lies in the interior of A. The same way we can treat the case when z lies in the interior of B.

Note that $\mathbb{E}^m$ admits an isometric splitting $\mathbb{E}^{m-2}\times\mathbb{E}^2$ such that

$$\dot{A}=\mathbb{E}^{m-2}\times\ddot{A}$$
$$\dot{B}=\mathbb{E}^{m-2}\times\ddot{B}$$

where $\ddot{A}$ and $\ddot{B}$ are half-planes in $\mathbb{E}^2$.

Using Exercise 3.7, let us replace each A by $\dot{A}$ and each B by $\dot{B}$ in the array, to get the array

$$(\underbrace{\dot{A},\dot{B},\dot{A},\dots}_{\lceil\frac{\pi}{\alpha}\rceil+1\text{ times}}).$$

The corresponding puff pastry $\dot{\mathcal{R}}$ splits as a product of $\mathbb{E}^{m-2}$ and a puff pastry, call it $\ddot{\mathcal{R}}$, glued from the copies of the plane $\mathbb{E}^2$ for the array

$$(\underbrace{\ddot{A}, \ddot{B}, \ddot{A}, \dots}_{\lceil\frac{\pi}{\alpha}\rceil + 1 \text{ times}}).$$

Note that the dihedral angle $\dot{A} \cap \dot{B}$ is at least α. Therefore the angle measure of $\ddot{A} \cap \ddot{B}$ is also at least α. According to Sublemma 3.6 and Observation 3.4, $\ddot{\mathcal{R}}$ is end-to-end convex.

Since $\dot{\mathcal{R}} \stackrel{iso}{=\!=} \mathbb{E}^{m-2} \times \ddot{\mathcal{R}}$, the puff pastry $\dot{\mathcal{R}}$ is also end-to-end convex; see 1.9.

It follows that the geodesic $[\dot{x}^0\dot{y}^N]_{\dot{\mathcal{R}}}$ contains $\dot{z}^0 = \dot{z}^N \in \dot{\mathcal{R}}$. By Proposition 3.2, the image of $[\dot{x}^0\dot{y}^N]_{\dot{\mathcal{R}}}$ under the map $\dot{x}^k \mapsto x^k$ is the geodesic $[x^0y^N]_{\mathcal{R}}$. Hence ❷ and the lemma follow. □

B Wide corners

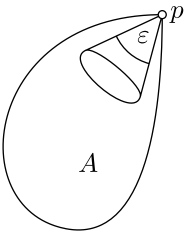

We say that a closed convex set $A \subset \mathbb{E}^m$ has ε-wide corners for given $\varepsilon > 0$ if together with each point p, the set A contains a small right circular cone with the tip at p and aperture ε; that is, ε is the maximum angle between two generating lines of the cone.

For example, a plane polygon has ε-wide corners if all its interior angles are at least ε.

We will consider arrays of closed convex sets $A^1, \dots, A^n$ in $\mathbb{E}^m$ such that for any subset $F \subset \{1, \dots, n\}$, the intersection $\bigcap_{i\in F} A^i$ has ε-wide corners. In this case, we may say briefly *all intersections of A^i have ε-wide corners.*

3.8. Exercise. *Assume $A^1, \dots, A^n \subset \mathbb{E}^m$ are compact, convex sets with a common interior point. Show that all intersections of A^i have ε-wide corners for some $\varepsilon > 0$.*

3.9. Exercise. *Assume $A^1, \dots, A^n \subset \mathbb{E}^m$ are convex sets with nonempty interiors that have a common center of symmetry. Show that all intersections of A^i have ε-wide corners for some $\varepsilon > 0$.*

The proof of the following proposition is based on 3.5; this lemma is essentially the case $n = 2$ in the proposition.

3.10. Proposition. *Given $\varepsilon > 0$ and a positive integer n, there is an array of integers $\boldsymbol{j}_\varepsilon(n) = (j_1, \dots, j_N)$ such that:*

(a) *For each k we have $1 \leqslant j_k \leqslant n$, and each number $1, \dots, n$ appears in $\boldsymbol{j}_\varepsilon$ at least once.*

(b) If $A^1, \dots, A^n$ is a collection of closed convex sets in $\mathbb{E}^m$ with a common point and all their intersections have ε-wide corners, then the puff pastry for the array $(A^{j_1}, \dots, A^{j_N})$ is end-to-end convex.

Moreover, we can assume that $N \leqslant (\lceil \frac{\pi}{\varepsilon} \rceil + 1)^n$.

Proof. The array $\boldsymbol{j}_\varepsilon(n) = (j_1, \dots, j_N)$ is constructed recursively. For $n = 1$, we can take $\boldsymbol{j}_\varepsilon(1) = (1)$.

Assume that $\boldsymbol{j}_\varepsilon(n)$ is constructed. Let us replace each occurrence of n in $\boldsymbol{j}_\varepsilon(n)$ by the alternating string

$$\underbrace{n, n+1, n, \dots}_{\lceil \frac{\pi}{\varepsilon} \rceil + 1 \text{ times}}.$$

Denote the obtained array by $\boldsymbol{j}_\varepsilon(n+1)$.

By Lemma 3.5, the end-to-end convexity of the puff pastry for $\boldsymbol{j}_\varepsilon(n+1)$ follows from the end-to-end convexity of the puff pastry for the array where each string

$$\underbrace{A^n, A^{n+1}, A^n, \dots}_{\lceil \frac{\pi}{\varepsilon} \rceil + 1 \text{ times}}$$

is replaced by $Q = A^n \cap A^{n+1}$. End-to-end convexity of the latter follows by the assumption on $\boldsymbol{j}_\varepsilon(n)$, since all the intersections of $A^1, \dots, A^{n-1}, Q$ have ε-wide corners.

The upper bound on N follows directly from the construction. $\square$

C Billiards

Let $A^1, A^2, \dots, A^n$ be a finite collection of closed convex sets in $\mathbb{E}^m$. Assume that for each i the boundary ∂A^i is a smooth hypersurface.

Consider the billiard table formed by the closure of the complement

$$T = \overline{\mathbb{E}^m \setminus \bigcup_i A^i}.$$

The sets A^i will be called walls of the table and the billiards described above will be called billiards with convex walls.

A billiard trajectory on the table is a unit-speed broken line γ that follows the standard law of billiards at the breakpoints on ∂A^i — in particular, the angle of reflection is equal to the angle of incidence. The breakpoints of the trajectory will be called collisions. *We assume that trajectory never meets a wall in a tangent direction. Also,* *we will always assume the trajectory meets at most one wall in a small neighborhood of every time moment. For example, if γ have to meet more then two walls at time moment t or collisions accumulate at t, then we have to assume that γ dies before t.*

Recall that the definition of sets with ε-wide corners is given in 3B.

3.11. Collision theorem. *Let $T \subset \mathbb{E}^m$ be a billiard table with n convex walls. Assume that the walls of T have a common interior point and all their intersections have ε-wide corners. Then the number of collisions of any trajectory in T is bounded by a number N which depends only on n and ε.*

As we will see from the proof, the value N can be found explicitly; $N = (\lceil\frac{\pi}{\varepsilon}\rceil + 1)^{n^2}$ will do.

3.12. Corollary. *Consider n homogeneous hard balls moving freely and colliding elastically in $\mathbb{R}^3$. Every ball moves along a straight line with constant speed until two balls collide, and then the new velocities of the two balls are determined by the laws of classical mechanics. We assume that only two balls can collide at the same time.*

Then the total number of collisions cannot exceed some number N that depends on the radii and masses of the balls. If the balls are identical, then N depends only on n.

The proof below admits a straightforward generalization to all dimensions.

3.13. Exercise. *Show that in the case of identical balls in the one-dimensional space (in $\mathbb{R}$) the total number of collisions cannot exceed $N = \frac{n\cdot(n-1)}{2}$.*

Proof of 3.12 modulo 3.11. Denote by $a_i = (x_i, y_i, z_i) \in \mathbb{R}^3$ the center of the i-th ball. Consider the corresponding point in $\mathbb{R}^{3\cdot N}$

$$\begin{aligned}\boldsymbol{a} &= (a_1, a_2, \dots, a_n) = \\ &= (x_1, y_1, z_1, x_2, y_2, z_2, \dots, x_n, y_n, z_n).\end{aligned}$$

The i-th and j-th balls intersect if

$$|a_i - a_j| \leqslant R_i + R_j,$$

where R_i denotes the radius of the i-th ball. These inequalities define $\frac{n\cdot(n-1)}{2}$ cylinders

$$C_{i,j} = \left\{ (a_1, a_2, \dots, a_n) \in \mathbb{R}^{3\cdot n} \;:\; |a_i - a_j| \leqslant R_i + R_j \right\}.$$

The closure of the complement

$$T = \overline{\mathbb{R}^{3\cdot n} \setminus \bigcup_{i<j} C_{i,j}}$$

is the configuration space of our system. Its points correspond to valid positions of the system of balls.

The evolution of the system of balls is described by the motion of the point $\boldsymbol{a} \in \mathbb{R}^{3\cdot n}$. It moves along a straight line at a constant speed until it hits one of the cylinders $C_{i,j}$; this event corresponds to a collision in the system of balls.

Consider the norm of $\boldsymbol{a} = (a_1, \dots, a_n) \in \mathbb{R}^{3\cdot n}$ defined by

$$\|\boldsymbol{a}\| = \sqrt{M_1{\cdot}|a_1|^2 + \dots + M_n{\cdot}|a_n|^2},$$

where $|a_i| = \sqrt{x_i^2 + y_i^2 + z_i^2}$ and M_i denotes the mass of the i-th ball. In the metric defined by $\|{*}\|$, the collisions follow the standard law of billiards.

By construction, the number of collisions of hard balls that we need to estimate is the same as the number of collisions of the corresponding billiard trajectory on the table with $C_{i,j}$ as the walls.

Note that each cylinder $C_{i,j}$ is a convex set; it has smooth boundary, and it is centrally symmetric around the origin. By 3.9, all the intersections of the walls have ε-wide corners for some $\varepsilon > 0$ that depend on the radiuses R_i and the masses M_i. It remains to apply the collision theorem (3.11). □

Now we present the proof of the collision theorem (3.11) based on the results developed in the previous section.

Proof of 3.11. Let us apply induction on n.

Base: $n = 1$. The number of collisions cannot exceed 1. Indeed, by the convexity of A^1, if the trajectory is reflected once in ∂A^1, then it cannot return to A^1.

Step. Assume γ is a trajectory that meets the walls in the order $A^{i_1}, \dots, A^{i_N}$ for a large integer N.

Consider the array

$$\boldsymbol{A}_\gamma = (A^{i_1}, \dots, A^{i_N}).$$

The induction hypothesis implies:

❶ *There is a positive integer M such that any M consecutive elements of $\boldsymbol{A}_\gamma$ contain each A^i at least once.*

Let $\mathcal{R}_\gamma$ be the Reshetnyak puff pastry for $\boldsymbol{A}_\gamma$.

Consider the lift of γ to $\mathcal{R}_\gamma$, defined by $\bar\gamma(t) = \gamma^k(t) \in \mathcal{R}_\gamma$ for any moment of time t between the k-th and $(k+1)$-th collisions. Since γ follows the standard law of billiards at breakpoints, the lift $\bar\gamma$ is locally a geodesic in $\mathcal{R}_\gamma$. By 3.1, the puff pastry $\mathcal{R}_\gamma$ is a proper geodesic CAT(0) space. Therefore $\bar\gamma$ is a geodesic.

Since γ does not meet $A^1 \cap \dots \cap A^n$, the lift $\bar\gamma$ does not lie in $\mathcal{R}^0_\gamma \cup \mathcal{R}^N_\gamma$. In particular, $\mathcal{R}_\gamma$ is not end-to-end convex.

Let

$$\boldsymbol{B} = (A^{j_1}, \dots, A^{j_K})$$

be the array provided by Proposition 3.10; so $\boldsymbol{B}$ contains each A^i at least once and the puff pastry $\mathcal{R}_{\boldsymbol{B}}$ for $\boldsymbol{B}$ is end-to-end convex. If N is sufficiently large, namely $N \geqslant K \cdot M$, then ❶ implies that $\boldsymbol{A}_\gamma$ can be obtained by inserting a finite number of A^i's in $\boldsymbol{B}$.

By 3.4, $\mathcal{R}_\gamma$ is end-to-end convex — a contradiction. □

3.14. Exercise. *Let $T \subset \mathbb{E}^m$ be a billiard table with n convex walls. Assume that the walls of T have a common point. Show that any trajectory in T has only finite number of collisions in a finite time interval.*

Construct an example of a billiard table on the plane with 2 convex walls such that the walls have a common point, but there is no upper bound on the number of collisions of its trajectories.

D Comments

The collision theorem (3.11) was proved by Dmitri Burago, Serge Ferleger and Alexey Kononenko [28]. Its corollary (3.12) answers a question posed by Yakov Sinai [47]. Puff pastry is used to bound topological entropy of the billiard flow and to approximate the shortest billiard path that touches given lines in a given order; see the papers of Dmitri Burago with Serge Ferleger, and Alexey Kononenko [29], and with Dimitri Grigoriev and Anatol Slissenko [30]. The lecture of Dmitri Burago [26] gives a short survey on the subject.

Note that the interior points of the walls play a key role in the proof despite that the trajectories never go inside the walls. In a similar fashion, puff pastry was used by Stephanie Alexander and Richard Bishop [4] to find the upper curvature bound for warped products.

Joel Hass [52] constructed an example of a Riemannian metric on the 3-ball with negative curvature and concave boundary. This example might decrease your appetite for generalizing the collision theorem — while locally such a 3-ball looks as good as the billiards table in the theorem, the number of collisions is obviously infinite.

It was shown by Dmitri Burago and Sergei Ivanov [31] that the number of collisions that may occur between n identical balls in $\mathbb{R}^3$ grows at least exponentially in n; the two-dimensional remains open.

Roman Barinov and Sergei Ivanov used another method to prove analogous statement in normed spaces with strongly convex smooth norm [57].

Lecture 4

Majorization

A Formulation

4.1. Definition. *Let $\mathcal{X}$ be a metric space, $\tilde{\alpha}$ be a simple closed curve of finite length in $\mathbb{E}^2$, and $D \subset \mathbb{E}^2$ be a closed region bounded by $\tilde{\alpha}$. A length-nonincreasing map $F\colon D \to \mathcal{X}$ is called m a j o r i z i n g if it is length-preserving on $\tilde{\alpha}$.*

In this case, we say that D m a j o r i z e s the curve $\alpha = F \circ \tilde{\alpha}$ under the map F.

The following proposition is a consequence of the definition.

4.2. Proposition. *Let α be a closed curve in a metric space $\mathcal{X}$. Suppose $D \subset \mathbb{E}^2$ majorizes α under $F\colon D \to \mathcal{X}$. Then any geodesic subarc of α is the image under F of a subarc of $\partial_{\mathbb{E}^2} D$ that is geodesic in the length metric of D.*

In particular, if D is convex, then the corresponding subarc is a geodesic in $\mathbb{E}^2$.

Proof. For a geodesic subarc $\gamma\colon [a,b] \to \mathcal{X}$ of $\alpha = F \circ \tilde{\alpha}$, set

$$\begin{aligned} \tilde{r} &= |\tilde{\gamma}(a) - \tilde{\gamma}(b)|_D, & \tilde{\gamma} &= (F|_{\partial D})^{-1} \circ \gamma, \\ s &= \operatorname{length} \gamma, & \tilde{s} &= \operatorname{length} \tilde{\gamma}. \end{aligned}$$

Then

$$\tilde{r} \geqslant r = s = \tilde{s} \geqslant \tilde{r}.$$

Therefore $\tilde{s} = \tilde{r}$. ☐

4.3. Corollary. *Assume a convex region $D \subset \mathbb{E}^2$ majorizes $[pxy]$. Then D is a s o l i d m o d e l t r i a n g l e of $[pxy]$; that is, $D = \operatorname{Conv}[\tilde{p}\tilde{x}\tilde{y}]$*

for a model triangle $[\tilde p\tilde x\tilde y]=\tilde\triangle(pxy)$. *Moreover, the majorizing map sends* $\tilde p$, $\tilde x$ *and* $\tilde y$ *respectively to* p, x *and* y.

Now we come to the main theorem of this section.

4.4. Majorization theorem. *Any closed rectifiable curve* α *in a geodesic* CAT(0) *space is majorized by a convex plane figure.*

This theorem is a useful tool for CAT(0) spaces. It can be used to get shorter proofs for several statements above; see for example 2.11.

B Triangles

The case when α is a triangle, say $[pxy]$, is the base in the following proof, and it is nontrivial. In this case, by Corollary 4.3, the majorizing convex region the solid model triangle.

4.5. Line-of-sight map. *Let* p *be a point and* α *be a curve of finite length in a geodesic space* $\mathcal{X}$. *Let* $\mathring\alpha:[0,1]\to\mathcal{U}$ *be the constant-speed parametrization of* α. *If* $\gamma_t\colon[0,1]\to\mathcal{U}$ *is a geodesic path from* p *to* $\mathring\alpha(t)$, *we say*

$$[0,1]\times[0,1]\to\mathcal{U}\colon(t,s)\mapsto\gamma_t(s)$$

is a line-of-sight map from p *to* α.

We will show that there is a majorizing map for $[pxy]$ whose image W is the image of the line-of-sight map for $[xy]$ from p. However, as one can see from the following example, the line-of-sight map is *not* majorizing in general.

Example. Let $\mathcal{Q}$ be a solid quadrangle $[pxzy]$ in $\mathbb{E}^2$ formed by two congruent triangles, which is non-convex at z (as in the picture). Equip $\mathcal{Q}$ with the length metric. Then $\mathcal{Q}$ is CAT(0) by Reshetnyak gluing (2.16). For triangle $[pxy]_{\mathcal{Q}}$ in $\mathcal{Q}$ and its model triangle $[\tilde p\tilde x\tilde y]$ in $\mathbb{E}^2$, we have

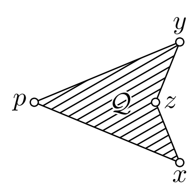

$$|\tilde x-\tilde y|=|x-y|_{\mathcal{Q}}=|x-z|+|z-y|.$$

Then the map F defined by matching line-of-sight parameters satisfies $F(\tilde x)=x$ and $|x-F(\tilde w)|>|\tilde x-\tilde w|$ if $\tilde w$ is near the midpoint $\tilde z$ of $[\tilde x\tilde y]$ and lies on $[\tilde p\tilde z]$. Indeed, for $\varepsilon=1-s$ we have

$$|\tilde x-\tilde w|=|\tilde x-\tilde\gamma_{\frac12}(s)|=|x-z|+o(\varepsilon)$$

and

$$|x-F(\tilde w)|=|x-\gamma_{\frac12}(s)|=|x-z|-\varepsilon\cdot\cos\measuredangle[z\,{}^{p}_{x}]+o(\varepsilon).$$

Thus F is not majorizing.

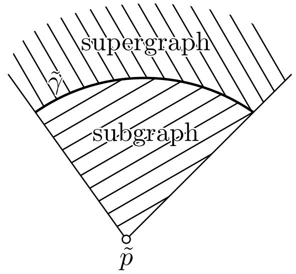

4.6. Definition. *Let $\tilde\gamma\colon \mathbb{I}\to\mathbb{E}^2$ be a curve and $\tilde p\in\mathbb{E}^2$ be such that the direction of $[\tilde p\,\tilde\gamma(t)]$ turns monotonically as t grows.*

The set formed by all geodesics from $\tilde p$ to the points on $\tilde\gamma$ is called the subgraph of $\tilde\gamma$ with respect to $\tilde p$.

The set of all points $\tilde x\in\mathbb{E}^2$ such that a geodesic $[\tilde p\tilde x]$ intersects $\tilde\gamma$ is called the supergraph of $\tilde\gamma$ with respect to $\tilde p$.

The curve $\tilde\gamma$ is called convex (concave) with respect to $\tilde p$ if the subgraph (supergraph) of $\tilde\gamma$ with respect to $\tilde p$ is convex.

This notion appears in [13] and an earlier form of it can be found in [63].

Our first lemma gives a model space construction based on repeated application of the argument in the proof of the inheritance lemma (2.15).

4.7. Lemma. *In $\mathbb{E}^2$, let β be a curve from x to y that is concave with respect to p. Let D be the subgraph of β with respect to p.*

(a) Then β forms a geodesic $[xy]_D$ in D and therefore β, $[px]$ and $[py]$ form a triangle $[pxy]_D$ in the length metric of D.

(b) Let $[\tilde p\tilde x\tilde y]$ be the model triangle for $[pxy]_D$. Then there is a short map

$$G\colon \operatorname{Conv}[\tilde p\tilde x\tilde y]\to D$$

such that $\tilde p\mapsto p$, $\tilde x\mapsto x$, $\tilde y\mapsto y$, and G is length-preserving on each side of $[\tilde p\tilde x\tilde y]$. In particular, $\operatorname{Conv}[\tilde p\tilde x\tilde y]$ *majorizes triangle $[pxy]_D$ in D under G.*

Proof. We prove the lemma for a polygonal line β; the general case then follows by approximation. Namely, since β is concave it can be approximated by polygonal lines that are concave with respect to p, with their lengths converging to $\operatorname{length}\beta$. Passing to a partial limit we will obtain the needed map G.

Suppose $\beta = x^0x^1\dots x^n$ is a polygonal line with $x^0=x$ and $x^n = y$. Consider a sequence of polygonal lines $\beta_i = x^0x^1\dots x^{i-1}y_i$ such that $|p-y_i|=|p-y|$ and β_i has same length as β; that is,

$$|x^{i-1}-y_i| = |x^{i-1}-x^i|+|x^i-x^{i+1}|+\dots+|x^{n-1}-x^n|.$$

Clearly $\beta_n=\beta$. Sequentially applying Alexandrov's lemma (1.12) shows that each of the polygonal lines $\beta_{n-1},\beta_{n-2},\dots,\beta_1$ is concave with respect to p. Let D_i be the subgraph of β_i with respect to p.

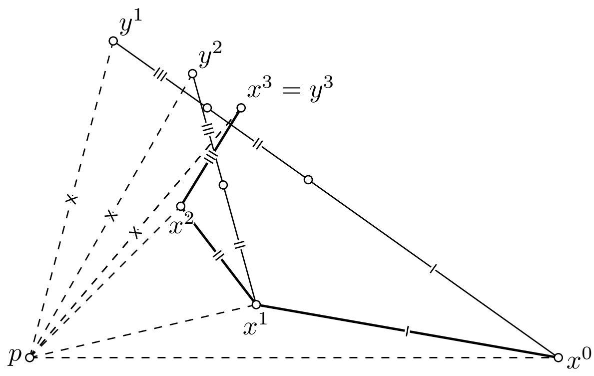

Applying the argument in the inheritance lemma (2.15) gives a short map $G_i\colon D_i \to D_{i+1}$ that maps $y_i \mapsto y_{i+1}$ and does not move p and x (in fact, G_i is the identity everywhere except on $\mathrm{Conv}[px^{i-1}y_i]$). Thus the composition

$$G_{n-1} \circ \dots \circ G_1\colon D_1 \to D_n$$

is short. The result follows since $D_1 \overset{iso}{=\!=} \mathrm{Conv}[\tilde{p}\tilde{x}\tilde{y}]$. □

4.8. Lemma. *Let $\mathcal{X}$ be a metric space, $\gamma\colon \mathbb{I} \to \mathcal{X}$ be a 1-Lipschitz curve, $p \in \mathcal{X}$, and $\tilde{p} \in \mathbb{E}^2$. Then there exists a unique up to rotation curve $\tilde{\gamma}\colon \mathbb{I} \to \mathbb{E}^2$, parametrized by arc-length, such that $|\tilde{p} - \tilde{\gamma}(t)| = |p - \gamma(t)|$ for all t and the direction of $[\tilde{p}\tilde{\gamma}(t)]$ monotonically turns around $\tilde{p}$ counterclockwise as t increases.*

If p, $\tilde{p}$, γ, and $\tilde{\gamma}$ are as above, then $\tilde{\gamma}$ is called the development of γ with respect to p; the point $\tilde{p}$ is called the basepoint of the development.

Proof. Consider the functions ρ, $\theta\colon \mathbb{I} \to \mathbb{R}$ defined as

$$\rho(t) = |p - \gamma(t)|, \qquad \theta(t) = \int\limits_{t_0}^{t} \frac{\sqrt{1-(\rho')^2}}{\rho},$$

where $t_0 \in \mathbb{I}$ is a fixed number and $\int$ denotes Lebesgue integral. Since γ is 1-Lipshitz, so is $\rho(t)$, and thus the function θ is defined and non-decreasing.

It is straightforward to check that (ρ, θ) uniquely describe $\tilde{\gamma}$ in polar coordinates on $\mathbb{E}^2$ with center at $\tilde{p}$. □

4.9. Exercise. *A geodesic space $\mathcal{U}$ is* CAT(0) *if and only if development of any geodesic with respect to any point is concave.*

4.10. Lemma. *Let $[pxy]$ be a triangle in a geodesic* CAT(0) *space $\mathcal{U}$. In $\mathbb{E}^2$, let $\tilde\gamma$ be the κ-development of $[xy]$ with respect to p, where $\tilde\gamma$ has basepoint $\tilde p$ and subgraph D. Consider the map $H\colon D\to\mathcal{U}$ that sends the point with parameter (t,s) under the line-of-sight map for $\tilde\gamma$ with respect to $\tilde p$, to the point with the same parameter under the line-of-sight map f for $[xy]$ with respect to p. Then H is length-nonincreasing. In particular, D majorizes triangle $[pxy]$.*

Proof. Let $\gamma\colon [0,T]\to\mathcal{U}$ be a unit-speed paremetrization of $[xy]$; so, $T=|x-y|$. Choose a partition

$$0=t^0<t^1<\dots<t^n=T,$$

and set $x^i=\gamma(t^i)$. Construct a chain of model triangles $[\tilde p\tilde x^{i-1}\tilde x^i]=$ $=\tilde\triangle(px^{i-1}x^i)$, with $\tilde x^0=\tilde x$ and the direction of $[\tilde p\tilde x^i]$ turning counterclockwise as i grows. Let D_n be the subgraph with respect to $\tilde p$ of the polygonal line $\tilde x^0\dots\tilde x^n$.

Let δ_n be the maximum radius of a circle inscribed in any of the triangles $[\tilde p\tilde x^{i-1}\tilde x^i]$.

Now we construct a map $H_n\colon D_n\to\mathcal{U}$ that increases distances by at most $2\cdot\delta_n$. Suppose $w\in D_n$. Then w lies on or inside some triangle $[\tilde p\tilde x^{i-1}\tilde x^i]$. Define $H_n(w)$ by first mapping w to a nearest point on $[\tilde p\tilde x^{i-1}\tilde x^i]$ (choosing one if there are several), followed by the natural map to the triangle $[px^{i-1}x^i]$.

Since triangles in $\mathcal{U}$ are thin, the restriction of H_n to each triangle $[\tilde p\tilde x^{i-1}\tilde x^i]$ is short. Then the triangle inequality implies that the restriction of H_n to

$$U_n=\bigcup_{1\leqslant i\leqslant n}[\tilde p\tilde x^{i-1}\tilde x^i]$$

is short with respect to the length metric on D_n. Since nearest-point projection from D_n to U_n increases the D_n-distance between two points by at most $2\cdot\delta_n$, the map H_n also increases the D_n-distance by at most $2\cdot\delta_n$.

Consider converging sequences $v_n\to v$ and $w_n\to w$ such that $v_n,w_n\in D_n$ and therefore $v,w\in D$. Note that

$$❶\qquad |H_n(v_n)-H_n(w_n)|\leqslant|v_n-w_n|_{D_n}+2\cdot\delta_n,$$

for each n. Since $\delta_n\to0$ and geodesics in $\mathcal{U}$ vary continuously with their endpoints (5.7), we have $H_n(v_n)\to H(v)$ and $H_n(w_n)\to H(w)$. Therefore the left-hand side in ❶ converges to $|H(v)-H(w)|$ and the right-hand side converges to $|v-w|_D$, it follows that H is short. $\square$

Proof of 4.4 for triangles. Suppose α is a triangle, say $[pxy]$.

Let $\tilde\gamma$ be the development of $[xy]$ with respect to p, where $\tilde\gamma$ has basepoint $\tilde p$ and subgraph D. By 4.9, $\tilde\gamma$ is concave. By 4.7, there is a short map $G\colon \operatorname{Conv}\tilde\triangle(pxy) \to D$. Further, by 4.10, D majorizes $[pxy]$ under a majorizing map $H\colon D \to \mathcal{U}$. Clearly $H \circ G$ is a majorizing map for $[pxy]$. $\square$

C Polygons

In the following proofs, $x^1 \dots x^n$ $(n \geqslant 3)$ denotes a polygonal line $x^1, \dots, x^n$, and $[x^1 \dots x^n]$ denotes the corresponding (closed) polygon. For a subset R of the ambient metric space, we denote by $[x^1 \dots x^n]_R$ a polygon in the length metric of R.

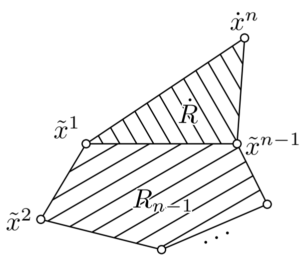

Proof of 4.4 for polygons. Now we claim that any closed n-gon $[x^1x^2 \dots x^n]$ in a CAT(0) space is majorized by a convex polygonal region

$$R_n = \operatorname{Conv}[\tilde x^1 \tilde x^2 \dots \tilde x^n]$$

under a map F_n such that $F_n\colon \tilde x^i \mapsto x^i$ for each i.

The base case $n = 3$ is proved above. Assume the statement is true for $(n-1)$-gons, $n \geqslant 4$. Then $[x^1x^2 \dots x^{n-1}]$ is majorized by a convex polygonal region

$$R_{n-1} = \operatorname{Conv}[\tilde x^1 \tilde x^2, \dots, \tilde x^{n-1}],$$

in $\mathbb{E}^2$ under a map F_{n-1} satisfying $F_{n-1}(\tilde x^i) = x^i$ for all i. Take $\dot x^n \in \mathbb{E}^2$ such that $[\tilde x^1 \tilde x^{n-1} \dot x^n] = \tilde\triangle(x^1 x^{n-1} x^n)$ and this triangle lies on the other side of $[\tilde x^1 \tilde x^{n-1}]$ from R_{n-1}. Let $\dot R = \operatorname{Conv}[\tilde x^1 \tilde x^{n-1} \dot x^n]$, and $\dot F\colon \dot R \to \mathcal{U}$ be a majorizing map for $[x^1 x^{n-1} x^n]$ as provided above.

Set $R = R_{n-1} \cup \dot R$, where R carries its length metric. Since F_n and F agree on $[\tilde x^1 \tilde x^{n-1}]$, we may define $F\colon R \to \mathcal{U}$ by

$$F(x) = \begin{cases} F_{n-1}(x), & x \in R_{n-1}, \\ \dot F(x), & x \in \dot R. \end{cases}$$

Then F is length-nonincreasing and is a majorizing map for $[x^1x^2 \dots$ $\dots x^n]$ (as in Definition 4.1).

If R is a convex subset of $\mathbb{E}^2$, we are done.

If R is not convex, the total internal angle of R at $\tilde x^1$ or $\tilde x^{n-1}$ or both is is larger than π. By relabeling we may suppose this holds for $\tilde x^{n-1}$.

The region R is obtained by gluing R_{n-1} to $\dot{R}$ by $[x^1x^{n-1}]$. Thus, by Reshetnyak gluing (2.16), R carrying its length metric is a CAT(0)-space. Moreover $[\tilde{x}^{n-2}\tilde{x}^{n-1}]\cup[\tilde{x}^{n-1}\dot{x}^n]$ is a geodesic of R. Thus $[\tilde{x}^1\tilde{x}^2\dots\tilde{x}^{n-2}\dot{x}^n]_R$ is a closed $(n-1)$-gon in R, to which the induction hypothesis applies. The resulting short map from a convex region in $\mathbb{E}^2$ to R, followed by F, is the desired majorizing map. □

For a polygon $[p_1\dots p_n]$, the values $\theta_i=\pi-\measuredangle[p_i\,{}^{p_{i-1}}_{p_{i+1}}]$ for all i (mod n) are called external angles of the polygon. The following exercise is a generalization of Fenchel's theorem.

4.11. Exercise. *Show that the sum of external angles of any polygon in a complete length* CAT(0) *space cannot be smaller than* $2\cdot\pi$.

The following exercise is a version of the Fáry–Milnor theorem for CAT(0) spaces.

4.12. Very advanced exercise. *Suppose that a simple polygon* β *in a complete length* CAT(0) *space does not bound an embedded disc. Show that the sum of external angles of* β *cannot be smaller than* $4\cdot\pi$.

Give an example of such a polygon β *with the sum of external angles exactly* $4\cdot\pi$.

4.13. Exercise. *Prove the following generalization of the arm lemma.*

4.14. Arm lemma. *Let* $P=[x^0x^1\dots x^{n+1}]$ *be a polygon in a geodesic* CAT(0) *space* $\mathcal{U}$. *Suppose* $\tilde{P}=[\tilde{x}^0\tilde{x}^1\dots\tilde{x}^{n+1}]$ *is a convex polygon in* $\mathbb{E}^2$ *such that*

❶ $$|\tilde{x}^i-\tilde{x}^{i-1}|_{\mathbb{E}^2}=|x^i-x^{i-1}|_{\mathcal{U}}\quad\text{and}\quad\measuredangle[x^i\,{}^{x^{i-1}}_{x^{i+1}}]\geqslant\measuredangle[\tilde{x}^i\,{}^{\tilde{x}^{i-1}}_{\tilde{x}^{i+1}}]$$

for all i. *Then* $|\tilde{x}^0-\tilde{x}^{n+1}|_{\mathbb{E}^2}\leqslant|x^0-x^{n+1}|_{\mathcal{U}}$.

D General case

If the space is proper, then the general case follows applying polygonal case to inscribed polygonal lines and passing to the limit.

E Comments

The last step in our proof essentially use that the space is proper. But the theorem holds for any geodesic CAT(0) space [9, 9.56].

The majorization theorem can be generalized to CAT(± 1) spaces; in the CAT(1) case one has to assume that the closed curve has length at most $2{\cdot}\pi$.

This theorem was proved by Yuriy Reshetnyak [79]; our proof uses a trick that we learned from the lectures of Werner Ballmann [19]. For complete spaces, another proof can be built on on the following closely related theorem. It was discovered by Urs Lang and Viktor Schroeder [61]; the third author proved it bit earlier, but did not publish the proof for quite a while [7, 9].

4.15. Generalized Kirszbraun's theorem. *Let $\mathcal{U}$ be a complete length* CAT(0) *space, let Q be arbitrary subset of a Euclidean space $\mathbb{E}^m$. Suppose $f\colon Q \to \mathcal{U}$ is a short map. Then $f\colon Q \to \mathcal{U}$ can be extended to a short map $F\colon \mathbb{E}^m \to \mathcal{U}$.*

4.16. Open problem. *Consider be a closed rectifiable curve α in a* CAT(0) *space $\mathcal{U}$. Note that if α is a geodesic triangle or it bounds an isometric copy of convex plane figure in $\mathcal{U}$, then α has a unique (up to congruence) majorizing convex figure.*

What about the converse?

Notice that the majorization theorem implies isoperimetrical inequality for CAT(0) metrics on the plane. Moreover, it implies the generalized isoperimetrical inequality in the style of Frederick Almgren [15]: *any curve of length $2{\cdot}\pi{\cdot}r$ in a* CAT(0) *length space spans a disc of area at most $\pi{\cdot}r^2$.*

The isoperimetrical inequality is known for 3- and 4-dimesnional Hadamard manifolds (that is, CAT(0) Riemannian manifolds). First the 4-dimensional case was proved by Christopher Croke [38], and latter Bruce Kleiner proved the 3-dimensional case [60]. Both papers are masterpieces. Despite many attempts, the isoperimetrical inequality for CAT(0) spaces remains open in all other dimensions and codimensions.

Lecture 5

Globalization

This lecture gives a sufficient condition for locally CAT(0) spaces to be globally CAT(0).

A Locally CAT spaces

We say that a space $\mathcal{U}$ is *locally* CAT(0) (or *locally* CAT(1)) if a small closed ball centered at any point p in $\mathcal{U}$ is CAT(0) (or CAT(1), respectively).

For example, the circle $\mathbb{S}^1 = \mathbb{R}/\mathbb{Z}$ is locally isometric to $\mathbb{R}$, and so $\mathbb{S}^1$ is locally CAT(0). On the other hand, $\mathbb{S}^1$ is not CAT(0), since closed local geodesics in $\mathbb{S}^1$ are not geodesics, so $\mathbb{S}^1$ does not meet 2.9.

If $\mathcal{U}$ is a proper geodesic space, then it is locally CAT(0) (or locally CAT(1)) if and only if each point $p \in \mathcal{U}$ admits an open neighborhood Ω that is geodesic and such that any triangle in Ω is thin (or spherically thin, respectively).

B Space of local geodesic paths

Recall that a constant-speed parameterization of a local geodesic by the unit interval $[0,1]$ is called a *local geodesic path*.

In this section, we will study the behavior of local geodesics in locally CAT(κ) spaces. The results will be used in the proof of the globalization theorem (5.6).

Recall that a *path* is a curve parametrized by $[0,1]$. The space of paths in a metric space $\mathcal{U}$ comes with the natural metric

❶ $$|\alpha - \beta| = \sup\{\, |\alpha(t) - \beta(t)|_{\mathcal{U}} \,:\, t \in [0,1]\,\}.$$

5.1. Proposition. *Let $\mathcal{U}$ be a proper geodesic, locally* CAT(κ) *space.*

Assume $\gamma_n\colon[0,1]\to\mathcal{U}$ is a sequence of local geodesic paths converging to a path $\gamma_\infty\colon[0,1]\to\mathcal{U}$. Then γ_∞ is a local geodesic path. Moreover

$$\operatorname{length}\gamma_n\to\operatorname{length}\gamma_\infty$$

as $n\to\infty$.

Proof; CAT(0) *case.* Fix $t\in[0,1]$. Let $R>0$ be sufficiently small, so that $\overline{\mathrm{B}}[\gamma_\infty(t),R]$ forms a proper geodesic CAT(0) space.

Assume that a local geodesic σ is shorter than $R/2$ and intersects the ball $\mathrm{B}(\gamma_\infty(t),R/2)$. Then σ cannot leave the ball $\overline{\mathrm{B}}[\gamma_\infty(t),R]$. By 2.9, σ is a geodesic. In particular, for all sufficiently large n, any arc of γ_n of length $R/2$ or less containing $\gamma_n(t)$ is a geodesic.

Since $\mathcal{B}=\overline{\mathrm{B}}[\gamma_\infty(t),R]$ is a proper geodesic CAT(0) space, by 2.2, geodesic segments in $\mathcal{B}$ depend uniquely on their endpoint pairs. Thus there is a subinterval $\mathbb{I}$ of $[0,1]$, that contains a neighborhood of t in $[0,1]$ and such that the arc $\gamma_n|_{\mathbb{I}}$ is minimizing for all large n. It follows that $\gamma_\infty|_{\mathbb{I}}$ is a geodesic, and therefore γ_∞ is a local geodesic.

The CAT(1) case is done in the same way, but one has to assume in addition that $R<\pi$. $\square$

The following lemma allows a local geodesic path to be moved continuously so that its endpoints follow given trajectories.

5.2. Patchwork along a geodesic. *Let $\mathcal{U}$ be a proper geodesic, locally* CAT(0) *space, and $\gamma\colon[0,1]\to\mathcal{U}$ be a locally geodesic path.*

Then there is a proper geodesic CAT(0) *space $\mathcal{N}$, an open set $\hat{\Omega}\subset\mathcal{N}$, and a geodesic path $\hat{\gamma}\colon[0,1]\to\hat{\Omega}$, such that there is an open locally distance-preserving map $\Phi\colon\hat{\Omega}\looparrowright\mathcal{U}$ satisfying $\Phi\circ\hat{\gamma}=\gamma$.*

If $\operatorname{length}\gamma<\pi$, *then the same holds in the* CAT(1) *case. Namely, we assume that $\mathcal{U}$ is a proper geodesic, locally* CAT(1) *space and construct a proper geodesic* CAT(1) *space $\mathcal{N}$ with the same property as above.*

Proof. Fix $r>0$ so that for each $t\in[0,1]$, the closed ball $\overline{\mathrm{B}}[\gamma(t),r]$ forms a proper geodesic CAT(0) space.

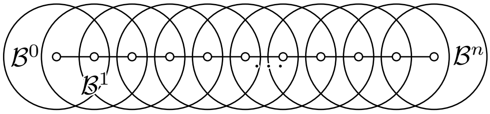

Choose a partition $0=t_0<t_1<\dots<t_n=1$ such that

$$\mathrm{B}(\gamma(t_i),r)\supset\gamma([t_{i-1},t_i])$$

for all $n > i > 0$. Set $\mathcal{B}_i = \overline{\mathrm{B}}[\gamma(t_i), r]$. We can assume in addition that $\mathcal{B}_{i-1} \cap \mathcal{B}_{i+1} \subset \mathcal{B}_i$ if $0 < i < n$.

Consider the disjoint union $\bigsqcup_i \mathcal{B}_i = \{\,(i, x) : x \in \mathcal{B}_i\,\}$ with the minimal equivalence relation $\sim$ such that $(i, x) \sim (i-1, x)$ for all i. Let $\mathcal{N}$ be the space obtained by gluing the $\mathcal{B}_i$ along $\sim$.

Note that $A_i = \mathcal{B}_i \cap \mathcal{B}_{i-1}$ is convex in $\mathcal{B}_i$ and in $\mathcal{B}_{i-1}$. Applying the Reshetnyak gluing theorem (2.16) n times, we conclude that $\mathcal{N}$ is a proper geodesic CAT(0) space.

For $t \in [t_{i-1}, t_i]$, define $\hat\gamma(t)$ as the equivalence class of $(i, \gamma(t))$ in $\mathcal{N}$. Let $\hat\Omega$ be the ε-neighborhood of $\hat\gamma$ in $\mathcal{N}$, where $\varepsilon > 0$ is chosen so that $\mathrm{B}(\gamma(t), \varepsilon) \subset \mathcal{B}_i$ for all $t \in [t_{i-1}, t_i]$.

Define $\Phi\colon \hat\Omega \to \mathcal{U}$ by sending the equivalence class of (i, x) to x. It is straightforward to check that Φ, $\hat\gamma$, and $\hat\Omega \subset \mathcal{N}$ satisfy the conclusion of the lemma.

The CAT(1) case is proved in the same way. $\square$

Recall that local geodesics are geodesics in any CAT(0) space; see 2.9. Using it with 5.2 and the uniqueness of geodesics (2.9), we get the following.

5.3. Corollary. *If $\mathcal{U}$ is a proper geodesic, locally* CAT(0) *space, then for any pair of points $p, q \in \mathcal{U}$, the space of all local geodesic paths from p to q is discrete; that is, for any local geodesic path γ connecting p to q, there is $\varepsilon > 0$ such that for any other local geodesic path δ from p to q we have $|\gamma(t) - \delta(t)|_{\mathcal{U}} > \varepsilon$ for some $t \in [0, 1]$.*

Analogously, if $\mathcal{U}$ is a proper geodesic, locally CAT(1) *space, then for any pair of points $p, q \in \mathcal{U}$, the space of all local geodesic paths shorter than π from p to q is discrete.*

5.4. Corollary. *If $\mathcal{U}$ is a proper geodesic, locally* CAT(0) *space, then for any path α there is a choice of local geodesic path γ_α connecting the ends of α such that the map $\alpha \mapsto \gamma_\alpha$ is continuous, and if α is a local geodesic path then $\gamma_\alpha = \alpha$.*

Analogously, if $\mathcal{U}$ is a proper geodesic, locally CAT(1) *space, then for any path α shorter than π, there is a choice of local geodesic path γ_α shorter than π connecting the ends of α such that the map $\alpha \mapsto \gamma_\alpha$ is continuous, and if α is a local geodesic path then $\gamma_\alpha = \alpha$.*

Proof of 5.4. We do the CAT(0) case; the CAT(1) case is analogous.

Consider the maximal interval $\mathbb{I} \subset [0, 1]$ containing 0 such that there is a continuous one-parameter family of local geodesic paths γ_t for $t \in \mathbb{I}$ connecting $\alpha(0)$ to $\alpha(t)$, with $\gamma_t(0) = \gamma_0(t) = \alpha(0)$ for any t.

By 5.1, $\mathbb{I}$ is closed, so we may assume $\mathbb{I} = [0, s]$ for some $s \in [0, 1]$.

Applying patchwork (5.2) to γ_s, we find that $\mathbb{I}$ is also open in $[0, 1]$. Hence $\mathbb{I} = [0, 1]$. Set $\gamma_\alpha = \gamma_1$.

By construction, if α is a local geodesic path, then $\gamma_\alpha = \alpha$.

Moreover, from 5.3, the construction $\alpha \mapsto \gamma_\alpha$ produces close results for sufficiently close paths in the metric defined by ❶; that is, the map $\alpha \mapsto \gamma_\alpha$ is continuous. $\square$

Given a path $\alpha\colon [0,1] \to \mathcal{U}$, we denote by $\bar\alpha$ the same path traveled in the opposite direction; that is,

$$\bar\alpha(t) = \alpha(1-t).$$

The p r o d u c t of two paths will be denoted with "$*$"; if two paths α and β connect the same pair of points, then the product $\bar\alpha * \beta$ is a closed curve.

5.5. Exercise. *Assume $\mathcal{U}$ is a proper geodesic, locally* CAT(1) *space. Consider the construction $\alpha \mapsto \gamma_\alpha$ provided by Corollary 5.4.*

Assume that α and β are two paths connecting the same pair of points in $\mathcal{U}$, where each is shorter than π and the product $\bar\alpha * \beta$ is null-homotopic in the class of closed curves shorter than $2\cdot\pi$. Show that $\gamma_\alpha = \gamma_\beta$.

C Globalization

5.6. Globalization theorem. *If a proper geodesic, locally* CAT(0) *space is simply connected, then it is* CAT(0).

Analogously, if $\mathcal{U}$ is a proper geodesic, locally CAT(1) *space such that any closed curve $\gamma\colon \mathbb{S}^1 \to \mathcal{U}$ shorter than $2\cdot\pi$ is null-homotopic in the class of closed curves shorter than $2\cdot\pi$. Then $\mathcal{U}$ is* CAT(1).

The surface on the diagram is an example of a simply connected space that is locally CAT(1) but not CAT(1). To contract the marked

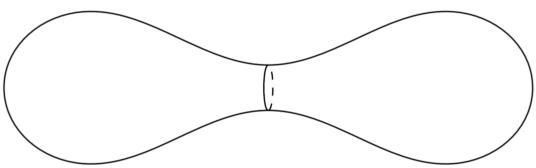

curve one has to increase its length to $2\cdot\pi$ or more; in particular, the surface does not satisfy the assumption of the globalization theorem.

The proof of the globalization theorem relies on the following theorem, which is essentially [10, Satz 9].

5.7. Patchwork globalization theorem. *A proper geodesic, locally* CAT(0) *space $\mathcal{U}$ is* CAT(0) *if and only if all pairs of points in $\mathcal{U}$*

are joined by unique geodesics, and these geodesics depend continuously on their endpoint pairs.

Analogously, a proper geodesic, locally CAT(1) *space* $\mathcal{U}$ *is* CAT(1) *if and only if all pairs of points in* $\mathcal{U}$ *at distance less than* π *are joined by unique geodesics, and these geodesics depend continuously on their endpoint pairs.*

The proof uses a thin-triangle decomposition with the inheritance lemma (2.15) and the line-of-sight map (4.5).

Proof of the patchwork globalization theorem (5.7). Note that the implication "only if" follows from 2.2 and 2.11; it remains to prove the "if" part.

Fix a triangle $[pxy]$ in $\mathcal{U}$. We need to show that $[pxy]$ is thin.

By the assumptions, the line-of-sight map $(t,s)\mapsto\gamma_t(s)$ from p to $[xy]$ is uniquely defined and continuous.

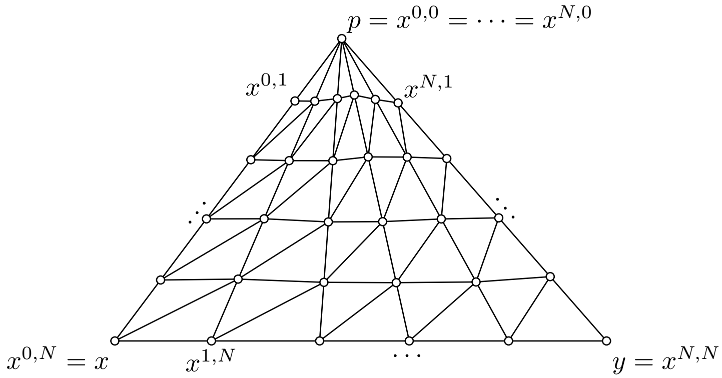

Fix a partition

$$0 = t^0 < t^1 < \ldots < t^N = 1,$$

and set $x^{i,j} = \gamma_{t^i}(t^j)$. Since the line-of-sight map is continuous and $\mathcal{U}$ is locally CAT(0), we may assume that the triangles

$$[x^{i,j}x^{i,j+1}x^{i+1,j+1}] \quad \text{and} \quad [x^{i,j}x^{i+1,j}x^{i+1,j+1}]$$

are thin for each pair i, j.

Now we show that the thin property propagates to $[pxy]$ by repeated application of the inheritance lemma (2.15):

◇ For fixed i, sequentially applying the lemma shows that the triangles $[px^{i,1}x^{i+1,2}]$, $[px^{i,2}x^{i+1,2}]$, $[px^{i,2}x^{i+1,3}]$, and so on are thin.

In particular, for each i, the long triangle $[px^{i,N}x^{i+1,N}]$ is thin.

◇ By the same lemma the triangles $[px^{0,N}x^{2,N}]$, $[px^{0,N}x^{3,N}]$, and so on, are thin.

In particular, $[pxy] = [px^{0,N}x^{N,N}]$ is thin. □

Proof of the globalization theorem; CAT(0) *case.* Let $\mathcal{U}$ be a proper geodesic, locally CAT(0) space that is simply connected. Given a path α in $\mathcal{U}$, denote by γ_α the local geodesic path provided by 5.4. Since the map $\alpha \mapsto \gamma_\alpha$ is continuous, by 5.3 we have $\gamma_\alpha = \gamma_\beta$ for any pair of paths α and β homotopic relative to the ends.

Since $\mathcal{U}$ is simply connected, any pair of paths with common ends are homotopic. In particular, if α and β are local geodesics from p to q, then $\alpha = \gamma_\alpha = \gamma_\beta = \beta$ by Corollary 5.4. It follows that any two points $p, q \in \mathcal{U}$ are joined by a unique local geodesic that depends continuously on (p, q).

Since $\mathcal{U}$ is geodesic, it remains to apply the patchwork globalization theorem (5.7).

CAT(1) *case.* The proof goes along the same lines, but one needs to use Exercise 5.5. □

5.8. Corollary. *Any compact geodesic, locally* CAT(0) *space that contains no closed local geodesics is* CAT(0)*.*

Analogously, any compact geodesic, locally CAT(1) *space that contains no closed local geodesics shorter than* $2{\cdot}\pi$ *is* CAT(1)*.*

Proof. By the globalization theorem (5.6), we need to show that the space is simply connected. Assume the contrary. Fix a nontrivial homotopy class of closed curves.

Denote by ℓ the exact lower bound for the lengths of curves in the class. Note that $\ell > 0$; otherwise, there would be a closed noncontractible curve in a CAT(0) neighborhood of some point, contradicting 2.4.

Since the space is compact, the class contains a length-minimizing curve, which must be a closed local geodesic.

The CAT(1) case is analogous, one only has to consider a homotopy class of closed curves shorter than $2{\cdot}\pi$. □

5.9. Exercise. *Prove that any compact geodesic, locally* CAT(0) *space* $\mathcal{X}$ *that is not* CAT(0) *contains a geodesic circle; that is, a simple closed curve* γ *such that for any two points* $p, q \in \gamma$*, one of the arcs of* γ *with endpoints* p *and* q *is a geodesic.*

Formulate and prove the analogous statement for CAT(1) *spaces.*

5.10. Advanced exercise. *Let* $\mathcal{U}$ *be a proper geodesic* CAT(0) *space. Assume* $\tilde{\mathcal{U}} \to \mathcal{U}$ *is a metric double cover branching along a geodesic*

γ. *(Formally speaking,* $\tilde{\mathcal{U}}$ *is completion of a double cover of the complement* $\mathcal{U} \setminus \gamma$*. For example, 3-dimensional Euclidean space admits a double cover branching along a line.)*

Show that $\tilde{\mathcal{U}}$ *is* CAT(0)*.*

D Remarks

The motivation for the notion of CAT(κ) spaces comes from the fact that a Riemannian manifold is locally CAT(κ) if and only if it has sectional curvature at most κ. This easily follows from Rauch comparison for Jacobi fields and Proposition 2.8.

The lemma about patchwork along a geodesic and its proof were suggested to us by Alexander Lytchak. This statement was originally proved by Stephanie Alexander and Richard Bishop [5] using a different method.

In the globalization theorem (5.6), properness can be weakened to completeness [see 9, and the references therein]. The original formulation of the globalization theorem, or Hadamard–Cartan theorem, states that *if M is a complete Riemannian manifold with sectional curvature at most* 0*, then the exponential map at any point $p \in M$ is a covering*; in particular, it implies that *the universal cover of M is diffeomorphic to the Euclidean space of the same dimension.*

In this generality, this theorem appeared in the lectures of Elie Cartan [34]. This theorem was proved for surfaces in Euclidean 3-space by Hans von Mangoldt [67] and a few years later independently for two-dimensional Riemannian manifolds by Jacques Hadamard [51].

Formulations for metric spaces of different generality were proved by Herbert Busemann [32], Willi Rinow [81], Mikhael Gromov [48, p. 119]. A detailed proof of Gromov's statement was given by Werner Ballmann [18] when $\mathcal{U}$ is proper, and by Stephanie Alexander and Richard Bishop [5] in more generality.

For proper CAT(1) spaces, the globalization theorem was proved by Brian Bowditch [23].

The globalization theorem holds for complete length spaces (not necessarily proper spaces) [9].

The patchwork globalization (5.7) is proved by Alexandrov [10, Satz 9]. For proper spaces one can remove the continuous dependence from the formulation; it follows from uniqueness. For complete spaces, the latter is not true [25, Chapter I, Exercise 3.14].

Lecture 6

Polyhedral spaces

This lecture describes a set of rules for gluing Euclidean cubes that produce a locally CAT(0) space and use these rules to construct exotic examples of aspherical manifolds.

A Products, cones, and suspension

Products, cones, and suspension are defined in Section 1D.

6.1. Proposition. *Let* $\mathcal{U}$ *and* $\mathcal{V}$ *be* CAT(0) *spaces. Then the product space* $\mathcal{U}\times\mathcal{V}$ *is* CAT(0)*.*

Proof. Fix a quadruple in $\mathcal{U}\times\mathcal{V}$:

$$p=(p_1,p_2),\qquad q=(q_1,q_2),\qquad x=(x_1,x_2),\qquad y=(y_1,y_2).$$

For the quadruple p_1,q_1,x_1,y_1 in $\mathcal{U}$, construct two model triangles $[\tilde p_1\tilde x_1\tilde y_1]=\tilde\triangle(p_1x_1y_1)_{\mathbb{E}^2}$ and $[\tilde q_1\tilde x_1\tilde y_1]=\tilde\triangle(q_1x_1y_1)_{\mathbb{E}^2}$. Similarly, for the quadruple p_2,q_2,x_2,y_2 in $\mathcal{V}$ construct two model triangles $[\tilde p_2\tilde x_2\tilde y_2]$ and $[\tilde q_2\tilde x_2\tilde y_2]$.

Consider four points in $\mathbb{E}^4=\mathbb{E}^2\times\mathbb{E}^2$

$$\tilde p=(\tilde p_1,\tilde p_2),\qquad \tilde q=(\tilde q_1,\tilde q_2),\qquad \tilde x=(\tilde x_1,\tilde x_2),\qquad \tilde y=(\tilde y_1,\tilde y_2).$$

Note that the triangles $[\tilde p\tilde x\tilde y]$ and $[\tilde q\tilde x\tilde y]$ in $\mathbb{E}^4$ are isometric to the model triangles $\tilde\triangle(pxy)_{\mathbb{E}^2}$ and $\tilde\triangle(qxy)_{\mathbb{E}^2}$.

If $\tilde z=(\tilde z_1,\tilde z_2)\in[\tilde x\tilde y]$, then $\tilde z_1\in[\tilde x_1\tilde y_1]$ and $\tilde z_2\in[\tilde x_2\tilde y_2]$ and

$$\begin{aligned}|\tilde z-\tilde p|^2_{\mathbb{E}^4}&=|\tilde z_1-\tilde p_1|^2_{\mathbb{E}^2}+|\tilde z_2-\tilde p_2|^2_{\mathbb{E}^2},\\ |\tilde z-\tilde q|^2_{\mathbb{E}^4}&=|\tilde z_1-\tilde q_1|^2_{\mathbb{E}^2}+|\tilde z_2-\tilde q_2|^2_{\mathbb{E}^2},\\ |p-q|^2_{\mathcal{U}\times\mathcal{V}}&=|p_1-q_1|^2_{\mathcal{U}}+|p_2-q_2|^2_{\mathcal{V}}.\end{aligned}$$

Therefore CAT(0) comparison for the quadruples p_1, q_1, x_1, y_1 in $\mathcal{U}$ and p_2, q_2, x_2, y_2 in $\mathcal{V}$ implies CAT(0) comparison for the quadruples p, q, x, y in $\mathcal{U} \times \mathcal{V}$. $\square$

Recall that metric on the cone c o n e $\mathcal{V} = \mathrm{Cone}\,\mathcal{U}$ is defined via cosine rule.

6.2. Proposition. *Let $\mathcal{U}$ be a metric space. Then* $\mathrm{Cone}\,\mathcal{U}$ *is* CAT(0) *if and only if $\mathcal{U}$ is* CAT(1).

Proof; if part. Given a point $x \in \mathrm{Cone}\,\mathcal{U}$, denote by x' its projection to $\mathcal{U}$ and by $|x|$ the distance from x to the tip of the cone; if x is the tip, then $|x| = 0$ and we can take any point of $\mathcal{U}$ as x'.

Let p, q, x, y be a quadruple in $\mathrm{Cone}\,\mathcal{U}$. Assume that the spherical model triangles $[\tilde{p}'\tilde{x}'\tilde{y}']_{\mathbb{S}^2} = \tilde{\triangle}(p'x'y')_{\mathbb{S}^2}$ and $[\tilde{q}'\tilde{x}'\tilde{y}']_{\mathbb{S}^2} = \tilde{\triangle}(q'x'y')_{\mathbb{S}^2}$ are defined. Consider the following points in $\mathbb{E}^3 = \mathrm{Cone}\,\mathbb{S}^2$:

$$\tilde{p} = |p|\cdot\tilde{p}', \qquad \tilde{q} = |q|\cdot\tilde{q}', \qquad \tilde{x} = |x|\cdot\tilde{x}', \qquad \tilde{y} = |y|\cdot\tilde{y}'.$$

Note that $[\tilde{p}\tilde{x}\tilde{y}]_{\mathbb{E}^3} \stackrel{iso}{=\!=} \tilde{\triangle}(pxy)_{\mathbb{E}^2}$ and $[\tilde{q}\tilde{x}\tilde{y}]_{\mathbb{E}^3} \stackrel{iso}{=\!=} \tilde{\triangle}(qxy)_{\mathbb{E}^2}$. Further, note that if $\tilde{z} \in [\tilde{x}\tilde{y}]_{\mathbb{E}^3}$, then $\tilde{z}' = \tilde{z}/|\tilde{z}|$ lies on the geodesic $[\tilde{x}'\tilde{y}']_{\mathbb{S}^2}$. Therefore the CAT(1) comparison for $|p' - q'|$ with $\tilde{z}' \in [\tilde{x}'\tilde{y}']_{\mathbb{S}^2}$ implies the CAT(0) comparison for $|p - q|$ with $\tilde{z} \in [\tilde{x}\tilde{y}]_{\mathbb{E}^3}$.

If at least one of the model triangles $\tilde{\triangle}(p'x'y')_{\mathbb{S}^2}$ and $\tilde{\triangle}(q'x'y')_{\mathbb{S}^2}$ is undefined, then the statement follows from the triangle inequalities

$$\begin{aligned} |p' - x'|_{\mathcal{U}} + |q' - x'|_{\mathcal{U}} &\geqslant |p' - q'|_{\mathcal{U}} \\ |p' - y'|_{\mathcal{U}} + |q' - y'|_{\mathcal{U}} &\geqslant |p' - q'|_{\mathcal{U}} \end{aligned}$$

This case is left as an exercise.

Only-if part. Suppose that $\tilde{p}'$, $\tilde{q}'$, $\tilde{x}'$, $\tilde{y}'$ are defined as above. Assume all these points lie in a half-space of $\mathbb{E}^3 = \mathrm{Cone}\,\mathbb{S}^2$ with origin at its boundary. Then we can choose positive values a, b, c, and d such that the points $a\cdot\tilde{p}'$, $b\cdot\tilde{q}'$, $c\cdot\tilde{x}'$, $d\cdot\tilde{y}'$ lie in one plane. Consider the corresponding points $a\cdot p'$, $b\cdot q'$, $c\cdot x'$, $d\cdot y'$ in $\mathrm{Cone}\,\mathcal{U}$. Applying the CAT(0) comparison for these points leads to CAT(1) comparison for the quadruple p', q', x', y' in $\mathcal{U}$.

It remains to consider the case when $\tilde{p}'$, $\tilde{q}'$, $\tilde{x}'$, $\tilde{y}'$ do not in a half-space. Fix $\tilde{z}' \in [\tilde{x}'\tilde{y}']_{\mathbb{S}^2}$. Observe that

$$|\tilde{p}' - \tilde{x}'|_{\mathbb{S}^2} + |\tilde{q}' - \tilde{x}'|_{\mathbb{S}^2} \leqslant |\tilde{p}' - \tilde{z}'|_{\mathbb{S}^2} + |\tilde{q}' - \tilde{z}'|_{\mathbb{S}^2}$$

or

$$|\tilde{p}' - \tilde{y}'|_{\mathbb{S}^2} + |\tilde{q}' - \tilde{y}'|_{\mathbb{S}^2} \leqslant |\tilde{p}' - \tilde{z}'|_{\mathbb{S}^2} + |\tilde{q}' - \tilde{z}'|_{\mathbb{S}^2}.$$

That is, in this case, the CAT(1) comparison follows from the triangle inequality. □

Suspension is a spherical analog of cone construction, and the following statement is a direct analog of 6.2; it can be proved along the same lines.

6.3. Proposition. *Let $\mathcal{U}$ be a metric space, and let $\mathcal{N}$ be a neighborhood of the north pole in* $\mathrm{Susp}\,\mathcal{U}$ *(possibly $\mathcal{N} = \mathrm{Susp}\,\mathcal{U}$) Then $\mathcal{N}$ is* CAT(1) *if and only if so is $\mathcal{U}$.*

B Polyhedral spaces

6.4. Definition. *A geodesic space $\mathcal{P}$ is called a (spherical) polyhedral space if it admits a finite triangulation τ such that every simplex in τ is isometric to a simplex in a Euclidean space (or respectively a unit sphere) of appropriate dimension.*

By triangulation of a polyhedral space, we will always understand a triangulation as above.

Note that according to the above definition, all polyhedral spaces are compact.

The dimension of a polyhedral space $\mathcal{P}$ is defined as the maximal dimension of the simplices in one (and therefore any) triangulation of $\mathcal{P}$.

Links. Let $\mathcal{P}$ be a polyhedral space and σ be a simplex in a triangulation τ of $\mathcal{P}$.

The simplices that contain σ form an abstract simplicial complex called the link of σ, denoted by Link_σ. If m is the dimension of σ, then the set of vertices of Link_σ is formed by the $(m+1)$-simplices that contain σ; the set of its edges is formed by the $(m+2)$-simplices that contain σ; and so on.

The link Link_σ can be identified with the subcomplex of τ formed by all the simplices σ' such that $\sigma \cap \sigma' = \varnothing$ but both σ and σ' are faces of a simplex of τ.

The points in Link_σ can be identified with the normal directions to σ at a point in its interior. The angle metric between directions makes Link_σ into a spherical polyhedral space. We will always consider the link with this metric.

Tangent space and space of directions. Let $\mathcal{P}$ be a polyhedral space (Euclidean or spherical) and τ be its triangulation. If a point

$p \in \mathcal{P}$ lies in the interior of a k-simplex σ of τ then the tangent space $\mathrm{T}_p = \mathrm{T}_p\mathcal{P}$ is naturally isometric to

$$\mathbb{E}^k \times (\operatorname{Cone}\operatorname{Link}_\sigma).$$

If $\mathcal{P}$ is an m-dimensional polyhedral space, then for any $p \in \mathcal{P}$ the space of directions Σ_p is a spherical polyhedral space of dimension at most $m-1$.

In particular, for any point p in σ, the isometry class of $\operatorname{Link}_\sigma$ together with $k = \dim\sigma$ determines the isometry class of Σ_p, and the other way around — Σ_p and k determines the isometry class of $\operatorname{Link}_\sigma$.

A small neighborhood of p is isometric to a neighborhood of the tip of $\operatorname{Cone}\Sigma_p$. In fact, if this property holds at any point of a compact length space $\mathcal{P}$, then $\mathcal{P}$ is a polyhedral space [62].

C CAT test

The following theorem provides a combinatorial description of polyhedral spaces with curvature bounded above.

6.5. Theorem. *Let τ be a triangulation of a polyhedral space $\mathcal{P}$. The space $\mathcal{P}$ is locally* CAT(0) *if and only if the link of each simplex in τ has no closed local geodesic shorter than $2\cdot\pi$.*

Analogously, let $\mathcal{P}$ be a spherical polyhedral space and τ be its triangulation. Then $\mathcal{P}$ is CAT(1) *if and only if neither $\mathcal{P}$ nor the link of any simplex in τ has a closed local geodesic shorter than $2\cdot\pi$.*

Proof. The "only if" part follows from 2.9, 6.3, and 6.2.

To prove the "if" part, we apply induction on $\dim\mathcal{P}$. The base case $\dim\mathcal{P} = 0$ is evident. Let us start with the CAT(1) case.

Step. Assume that the theorem is proved in the case $\dim\mathcal{P} < m$. Suppose $\dim\mathcal{P} = m$.

Fix a point $p \in \mathcal{P}$. A neighborhood of p is isometric to a neighborhood of the north pole in the suspension over the space of directions Σ_p.

Note that Σ_p is a spherical polyhedral space, and its links are isometric to links of $\mathcal{P}$. By the induction hypothesis, Σ_p is CAT(1). Thus, by the second part of Exercise 6.2, $\mathcal{P}$ is locally CAT(1).

Applying the second part of Corollary 5.8, we get the statement.

The CAT(0) case is done in exactly the same way except we need to use Proposition 6.2 and the first part of Corollary 5.8 on the last step. □

6.6. Exercise. *Let $\mathcal{P}$ be a polyhedral space such that any two points can be connected by a unique geodesic. Show that $\mathcal{P}$ is* CAT(0).

6.7. Advanced exercise. *Construct a Euclidean polyhedral metric on $\mathbb{S}^3$ such that the total angle around each edge in its triangulation is at least $2\cdot\pi$.*

D Flag complexes

6.8. Definition. *A simplicial complex $\mathcal{S}$ is called flag if whenever $\{v^0,\dots,v^k\}$ is a set of distinct vertices of $\mathcal{S}$ that are pairwise joined by edges, then the vertices $v^0,\dots,v^k$ span a k-simplex in $\mathcal{S}$.*

If the above condition is satisfied for $k = 2$, then we say that $\mathcal{S}$ satisfies the no-triangle condition.

Note that every flag complex is determined by its one-skeleton. Moreover, for any graph, its cliques (that is, complete subgraphs) define a flag complex. For that reason, flag complexes are also called clique complexes.

6.9. Exercise. *Show that the barycentric subdivision of any simplicial complex is a flag complex.*

Use the flag condition (see 6.12 below) to conclude that any finite simplicial complex is homeomorphic to a proper length $\mathrm{CAT}(1)$ *space.*

6.10. Proposition. *A simplicial complex $\mathcal{S}$ is flag if and only if $\mathcal{S}$ as well as the links of all its simplices satisfy the no-triangle condition.*

From the definition of flag complex, we get the following.

6.11. Observation. *Any link of any simplex in a flag complex is flag.*

Proof of 6.10. By Observation 6.11, the no-triangle condition holds for any flag complex and the links of all its simplices.

Now assume that a complex $\mathcal{S}$ and all its links satisfy the no-triangle condition. It follows that $\mathcal{S}$ includes a 2-simplex for each triangle. Applying the same observation for each edge we get that $\mathcal{S}$ includes a 3-simplex for any complete graph with 4 vertices. Repeating this observation for triangles, 4-simplices, 5-simplices, and so on, we get that $\mathcal{S}$ is flag. □

All-right triangulation. A triangulation of a spherical polyhedral space is called an all-right triangulation if each simplex of the triangulation is isometric to a spherical simplex all of whose angles are right. Similarly, we say that a simplicial complex is equipped with an all-right spherical metric if it is a length metric and each simplex is isometric to a spherical simplex all of whose angles are right.

Spherical polyhedral CAT(1) spaces glued from right-angled simplices admit the following characterization discovered by Mikhael Gromov [48, p. 122].

6.12. Flag condition. *Assume that a spherical polyhedral space $\mathcal{P}$ admits an all-right triangulation τ. Then $\mathcal{P}$ is* CAT(1) *if and only if τ is flag.*

Proof; only-if part. Assume there are three vertices v^1, v^2, and v^3 of τ that are pairwise joined by edges but do not span a triangle. Note that in this case

$$\measuredangle[v^1\,{}^{v^2}_{v^3}] = \measuredangle[v^2\,{}^{v^3}_{v^1}] = \measuredangle[v^3\,{}^{v^1}_{v^2}] = \pi.$$

Equivalently,

❶ *The product of the geodesics $[v^1v^2]$, $[v^2v^3]$, and $[v^3v^1]$ forms a locally geodesic loop in $\mathcal{P}$ of length $\tfrac{3}{2}\cdot\pi$.*

Now assume that $\mathcal{P}$ is CAT(1). Then by 6.3, $\mathrm{Link}_\sigma\mathcal{P}$ is CAT(1) for every simplex σ in τ.

Each of these links is an all-right spherical complex and by 5.8, none of these links can contain a geodesic circle shorter than $2\cdot\pi$.

Therefore Proposition 6.10 and ❶ imply the "only if" part.

If part. By 6.11 and 5.8, it is sufficient to show that any closed local geodesic γ in a flag complex $\mathcal{S}$ with all-right metric has length at least $2\cdot\pi$.

Recall that the closed star of a vertex v (briefly $\overline{\mathrm{Star}}_v$) is formed by all the simplices containing v. Similarly, Star_v, the open star of v, is the union of all simplices containing v with faces opposite v removed.

Choose a vertex v such that Star_v contains a point $\gamma(t_0)$ of γ. Consider the maximal arc γ_v of γ that contains the point $\gamma(t_0)$ and runs in Star_v. Note that the distance $|v-\gamma_v(t)|_{\mathcal{P}}$ behaves in exactly the same way as the distance from the north pole in $\mathbb{S}^2$ to a geodesic in the northern hemisphere; that is, there is a geodesic $\tilde\gamma_v$ in the northern hemisphere of $\mathbb{S}^2$ such that for any t we have

$$|v-\gamma_v(t)|_{\mathcal{P}} = |n-\tilde\gamma_v(t)|_{\mathbb{S}^2},$$

where n denotes the north pole of $\mathbb{S}^2$. In particular,

$$\mathrm{length}\,\gamma_v = \pi;$$

that is, γ spends time π on every visit to Star_v.

After leaving Star_v, the local geodesic γ has to enter another simplex, say σ'. Since τ is flag, the simplex σ' has a vertex v' not joined to v by an edge; that is,

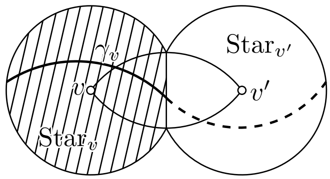

$$\mathrm{Star}_v \cap \mathrm{Star}_{v'} = \varnothing$$

The same argument as above shows that γ spends time π on every visit to $\mathrm{Star}_{v'}$. Therefore the total length of γ is at least $2{\cdot}\pi$. $\square$

6.13. Exercise. *Assume that a spherical polyhedral space $\mathcal{P}$ admits a triangulation τ such that all edge lengths of all simplices are at least $\frac{\pi}{2}$. Show that $\mathcal{P}$ is* CAT(1) *if τ is flag.*

6.14. Exercise. *Let P be a convex polyhedron in $\mathbb{E}^3$ with n faces $F_1, \dots, F_n$. Suppose that each face of P has only obtuse or right angles. Let us take 2^n copies of P indexed by an n-bit array. Glue two copies of P along F_i if their arrays differ only in the i-th bit. Show that the obtained space is a locally* CAT(0) *topological manifold.*

The space of trees. The following construction is given by Louis Billera, Susan Holmes, and Karen Vogtmann [20].

Let $\mathcal{T}_n$ be the set of all metric trees with n end vertices labeled by $a^1, \dots, a^n$. To describe one tree in $\mathcal{T}_n$ we may fix a topological tree t with end vertices $a^1, \dots, a^n$, and all other vertices of degree 3, and prescribe the lengths of $2{\cdot}n - 3$ edges. If the length of an edge vanishes, we assume that this edge degenerates; such a tree can be also described using a different topological tree t'. The subset of $\mathcal{T}_n$ corresponding to the given topological tree t can be identified with the octant

$$\left\{\, (x_1, \dots, x_{2\cdot n-3}) \in \mathbb{R}^{2\cdot n-3} \;:\; x_i \geqslant 0 \,\right\}.$$

Equip each such subset with the metric induced from $\mathbb{R}^{2\cdot n-3}$ and consider the length metric on $\mathcal{T}_n$ induced by these metrics.

6.15. Exercise. *Show that $\mathcal{T}_n$ with the described metric is* CAT(0)*.*

E Cubical complexes

The definition of a cubical complex mostly repeats the definition of a simplicial complex, with simplices replaced by cubes.

Formally, a *cubical complex* is defined as a subcomplex of the unit cube in the Euclidean space $\mathbb{R}^N$ of large dimension; that is, a collection of faces of the cube such that together with each face it contains all its sub-faces. Each cube face in this collection will be called a cube of the cubical complex.

Note that according to this definition, any cubical complex is finite.

The union of all the cubes in a cubical complex $\mathcal{Q}$ will be called its underlying space. A homeomorphism from the underlying space of $\mathcal{Q}$ to a topological space $\mathcal{X}$ is called a cubulation of $\mathcal{X}$.

The underlying space of a cubical complex $\mathcal{Q}$ will be always considered with the length metric induced from $\mathbb{R}^N$. In particular, with this metric, each cube of $\mathcal{Q}$ is isometric to the unit cube of the corresponding dimension.

It is straightforward to construct a triangulation of the underlying space of $\mathcal{Q}$ such that each simplex is isometric to a Euclidean simplex. In particular, the underlying space of $\mathcal{Q}$ is a Euclidean polyhedral space.

The link of a cube in a cubical complex is defined similarly to the link of a simplex in a simplicial complex. It is a simplicial complex that admits a natural all-right triangulation — each simplex corresponds to an adjusted cube.

Cubical analog of a simplicial complex. Let $\mathcal{S}$ be a finite simplicial complex and $\{v_1,\dots,v_N\}$ be the set of its vertices.

Consider $\mathbb{R}^N$ with the standard basis $\{e_1,\dots,e_N\}$. Denote by $\square^N$ the standard unit cube in $\mathbb{R}^N$; that is,

$$\square^N = \left\{\, (x_1,\dots,x_N)\in\mathbb{R}^N \;:\; 0\leqslant x_i\leqslant 1 \text{ for each } i \,\right\}.$$

Given a k-dimensional simplex $\langle v_{i_0},\dots,v_{i_k}\rangle$ in $\mathcal{S}$, mark the $(k+1)$-dimensional faces in $\square^N$ (there are 2^{N-k} of them) which are parallel to the coordinate $(k+1)$-plane spanned by $e_{i_0},\dots,e_{i_k}$.

Note that the set of all marked faces of $\square^N$ forms a cubical complex; it will be called the cubical analog of $\mathcal{S}$ and will be denoted as $\square_{\mathcal{S}}$.

6.16. Proposition. *Let $\mathcal{S}$ be a finite connected simplicial complex and $\mathcal{Q}=\square_{\mathcal{S}}$ be its cubical analog. Then the underlying space of $\mathcal{Q}$ is connected and the link of any vertex of $\mathcal{Q}$ is isometric to $\mathcal{S}$ equipped with the all-right spherical metric.*

In particular, if $\mathcal{S}$ is a flag complex, then $\mathcal{Q}$ is a locally CAT(0)*, and therefore its universal cover $\tilde{\mathcal{Q}}$ is* CAT(0).

Proof. The first part of the proposition follows from the construction of $\square_{\mathcal{S}}$.

If $\mathcal{S}$ is flag, then by the flag condition (6.12) the link of any cube in $\mathcal{Q}$ is CAT(1). Therefore, by the cone construction (6.2) $\mathcal{Q}$ is locally CAT(0). It remains to apply the globalization theorem (5.6). $\square$

From Proposition 6.16, it follows that the cubical analog of any flag complex is aspherical. The following exercise states that the converse also holds; see [41, 5.4].

6.17. Exercise. *Show that a finite simplicial complex is flag if and only if its cubical analog is aspherical.*

F Construction

By 2.4, any complete length CAT(0) space is contractible. Therefore, by the globalization theorem (5.6), all proper length, locally CAT(0) spaces are aspherical; that is, they have contractible universal covers. This observation will be used to construct examples of aspherical spaces.

Let $\mathcal{X}$ be a proper topological space. Recall that $\mathcal{X}$ is called simply connected at infinity if for any compact set $K \subset \mathcal{X}$ there is a bigger compact set $K' \supset K$ such that $\mathcal{X} \setminus K'$ is path-connected and any loop which lies in $\mathcal{X} \setminus K'$ is null-homotopic in $\mathcal{X} \setminus K$.

Recall that path-connected spaces are not empty by definition. Therefore compact spaces are not simply connected at infinity.

The following example was constructed by Michael Davis [40].

6.18. Proposition. *For any $m \geqslant 4$, there is a closed aspherical m-dimensional manifold whose universal cover is not simply connected at infinity.*

In particular, the universal cover of this manifold is not homeomorphic to the m-dimensional Euclidean space.

The proof requires the following lemma.

6.19. Lemma. *Let $\mathcal{S}$ be a finite flag complex, $\mathcal{Q} = \square_{\mathcal{S}}$ be its cubical analog and $\tilde{\mathcal{Q}}$ be the universal cover of $\mathcal{Q}$.*

Assume $\tilde{\mathcal{Q}}$ is simply connected at infinity. Then $\mathcal{S}$ is simply connected.

Proof. Assume $\mathcal{S}$ is not simply connected. Equip $\mathcal{S}$ with an all-right spherical metric. Choose a shortest noncontractible circle $\gamma\colon \mathbb{S}^1 \to \mathcal{S}$ formed by the edges of $\mathcal{S}$.

Note that γ forms a one-dimensional subcomplex of $\mathcal{S}$ which is a closed local geodesic. Denote by G the subcomplex of $\mathcal{Q}$ which corresponds to γ.

Fix a vertex $v \in G$; let G_v be the connected component of v in G. Let $\tilde{G}$ be a connected component of the inverse image of G_v in $\tilde{\mathcal{Q}}$ for the universal cover $\tilde{\mathcal{Q}} \to \mathcal{Q}$. Fix a point $\tilde{v} \in \tilde{G}$ in the inverse image of v.

Note that

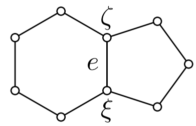

❶ *$\tilde{G}$ is a convex set in $\tilde{\mathcal{Q}}$.*

Indeed, according to Proposition 6.16, $\tilde{\mathcal{Q}}$ is CAT(0). By Exercise 2.14, it is sufficient to show that $\tilde{G}$ is locally convex in $\tilde{\mathcal{Q}}$, or equivalently, G is locally convex in $\mathcal{Q}$.

Note that the latter can only fail if γ contains two vertices, say ξ and ζ in $\mathcal{S}$, which are joined by an edge not in γ; denote this edge by e.

Each edge of $\mathcal{S}$ has length $\frac{\pi}{2}$. Therefore each of the two circles formed by e and an arc of γ from ξ to ζ is shorter than γ. Moreover, at least one of them is noncontractible since γ is noncontractible. That is, γ is not a shortest noncontractible circle, a contradiction. $\triangle$

Further, note that $\tilde{G}$ is homeomorphic to the plane since $\tilde{G}$ is a two-dimensional manifold without boundary which by the above is CAT(0) and hence is contractible.

Denote by C_R the circle of radius R in $\tilde{G}$ centered at $\tilde{v}$. All C_R are homotopic to each other in $\tilde{G}\setminus\{\tilde{v}\}$ and therefore in $\tilde{\mathcal{Q}}\setminus\{\tilde{v}\}$.

Note that the map $\tilde{\mathcal{Q}}\setminus\{\tilde{v}\}\to\mathcal{S}$ which returns the direction of $[\tilde{v}x]$ for any $x\neq\tilde{v}$, maps C_R to a circle homotopic to γ. Therefore C_R is not contractible in $\tilde{\mathcal{Q}}\setminus\{\tilde{v}\}$.

If R is large, the circle C_R lies outside of any fixed compact set K' in $\tilde{\mathcal{Q}}$. From above C_R is not contractible in $\tilde{\mathcal{Q}}\setminus K$ if $K\supset\tilde{v}$. It follows that $\tilde{\mathcal{Q}}$ is not simply connected at infinity, a contradiction. $\square$

The proof of the following exercise is analogous. It will be used later in the proof of Proposition 6.21 — a more geometric version of Proposition 6.18.

6.20. Exercise. *Under the assumptions of Lemma 6.19, for any vertex v in $\mathcal{S}$ the complement $\mathcal{S}\setminus\{v\}$ is simply connected.*

Proof of 6.18. Let Σ^{m-1} be an $(m-1)$-dimensional smooth homology sphere that is not simply connected, and bounds a contractible smooth compact m-dimensional manifold $\mathcal{W}$.

For $m\geqslant 5$, the existence of such $(\mathcal{W},\Sigma)$ is proved by Michel Kervaire [58]. For $m=4$, it follows from the construction of Barry Mazur [68].

Pick any triangulation τ of W and let $\mathcal{S}$ be the resulting subcomplex that triangulates Σ.

We can assume that $\mathcal{S}$ is flag; otherwise, pass to the barycentric subdivision of τ and apply Exercise 6.9.

Let $\mathcal{Q}=\square_{\mathcal{S}}$ be the cubical analog of $\mathcal{S}$.

By Proposition 6.16, $\mathcal{Q}$ is a homology manifold. It follows that $\mathcal{Q}$ is a piecewise linear manifold with a finite number of singularities at its vertices.

Removing a small contractible neighborhood V_v of each vertex v in $\mathcal{Q}$, we can obtain a piecewise linear manifold $\mathcal{N}$ whose boundary is formed by several copies of Σ.

Let us glue a copy of $\mathcal{W}$ along its boundary to each copy of Σ in the boundary of $\mathcal{N}$. This results in a closed manifold $\mathcal{M}$ with polyhedral metric which is homotopically equivalent to $\mathcal{Q}$.

Indeed, since both V_v and $\mathcal{W}$ are contractible, the identity map of their common boundary Σ can be extended to a homotopy equivalence $V_v \to \mathcal{W}$ relative to the boundary. Therefore the identity map on $\mathcal{N}$ extends to homotopy equivalences $f\colon \mathcal{Q} \to \mathcal{M}$ and $g\colon \mathcal{M} \to \mathcal{Q}$.

Finally, by Lemma 6.19, the universal cover $\tilde{\mathcal{Q}}$ of $\mathcal{Q}$ is not simply connected at infinity.

The same holds for the universal cover $\tilde{\mathcal{M}}$ of $\mathcal{M}$. The latter follows since the constructed homotopy equivalences $f\colon \mathcal{Q} \to \mathcal{M}$ and $g\colon \mathcal{M} \to \to \mathcal{Q}$ lift to *proper maps* $\tilde{f}\colon \tilde{\mathcal{Q}} \to \tilde{\mathcal{M}}$ and $\tilde{g}\colon \tilde{\mathcal{M}} \to \tilde{\mathcal{Q}}$; that is, for any compact sets $A \subset \tilde{\mathcal{Q}}$ and $B \subset \tilde{\mathcal{M}}$, the inverse images $\tilde{g}^{-1}(A)$ and $\tilde{f}^{-1}(B)$ are compact. □

The following proposition was proved by Fredric Ancel, Michael Davis, and Craig Guilbault [16]; it could be considered as a more geometric version of Proposition 6.18.

6.21. Proposition. *Given $m \geqslant 5$, there is a Euclidean polyhedral space $\mathcal{P}$ such that:*

(a) $\mathcal{P}$ is homeomorphic to a closed m-dimensional manifold.

(b) $\mathcal{P}$ is locally $\mathrm{CAT}(0)$.

(c) The universal cover of $\mathcal{P}$ is not simply connected at infinity.

Dale Rolfsen [82] has shown that there are no three-dimensional examples of that type. Paul Thurston [87] conjectured that the same holds in the four-dimensional case.

Proof. Apply Exercise 6.20 to the barycentric subdivision of the simplicial complex $\mathcal{S}$ provided by Exercise 6.22. □

6.22. Exercise. *Given an integer $m \geqslant 5$, construct a finite $(m-1)$-dimensional simplicial complex $\mathcal{S}$ such that* $\mathrm{Cone}\,\mathcal{S}$ *is homeomorphic to $\mathbb{E}^m$ and $\pi_1(\mathcal{S} \setminus \{v\}) \neq 0$ for some vertex v in $\mathcal{S}$.*

G Remarks

Theorem 6.5 gives is a good-looking description of polyhedral $\mathrm{CAT}(\kappa)$ spaces, but in fact, it is hard to check even in very simple cases. For example, the description of those coverings of $\mathbb{S}^3$ branching at three great circles which are CAT(1) requires quite a bit of work [35] — try to guess the answer before reading.

Another example is the braid space $\mathcal{B}_n$ that is the universal cover of $\mathbb{C}^n$ infinitely branching in complex hyperplanes $z_i = z_j$ with the induced length metric. So far it is not known if $\mathcal{B}_n$ is CAT(0) for any $n \geqslant 4$ [73]. Understanding this space could help to study the braid group. This circle of questions is closely related to the generalization of the flag condition (6.12) to spherical simplices with few acute dihedral angles.

The construction used in the proof of Proposition 6.18 admits a number of modifications, several of which are discussed in a survey by Michael Davis [41].

A similar argument was used by Michael Davis, Tadeusz Januszkiewicz, and Jean-François Lafont [43]. They constructed a closed smooth four-dimensional manifold M with universal cover $\tilde{M}$ diffeomorphic to $\mathbb{R}^4$, such that M admits a polyhedral metric which is locally CAT(0), but does not admit a Riemannian metric with nonpositive sectional curvature. Another example of that type was constructed by Stephan Stadler [85]. There are no lower-dimensional examples of this type — the two-dimensional case follows from the classification of surfaces, and the three-dimensional case follows from the geometrization conjecture.

It is noteworthy that any complete, simply connected Riemannian manifold with nonpositive curvature is homeomorphic to the Euclidean space of the same dimension. In fact, by the globalization theorem (5.6), the exponential map at a point of such a manifold is a homeomorphism. In particular, there is no Riemannian analog of Proposition 6.21.

Recall that a triangulation of an m-dimensional manifold defines a piecewise linear structure if the link of every simplex Δ is homeomorphic to the sphere of dimension $m - 1 - \dim\Delta$. According to Stone's theorem [42, 86], the triangulation of $\mathcal{P}$ in Proposition 6.21 cannot be made piecewise linear — despite the fact that $\mathcal{P}$ is a manifold, its triangulation does not induce a piecewise linear structure.

The flag condition also leads to the so-called hyperbolization procedure, a flexible tool for constructing aspherical spaces; a good survey on the subject is given by Ruth Charney and Michael Davis [36].

The CAT(0) property of a cube complex admits interesting (and useful) geometric descriptions if one exchanged the ℓ^2-metric to a natural ℓ^1 or ℓ^∞ on each cube.

6.23. Theorem. *The following three conditions are equivalent.*

(a) A cube complex Q equipped with ℓ^2-metric is CAT(0).

(b) A cube complex Q equipped with ℓ^∞-metric is injective.

(c) A cube complex Q equipped with ℓ^1-metric is median. The latter means that for any three points x, y, z there is a unique point m (it is called the median of x, y, and z) and a choice of geodesics such that $[xy] \ni m$, $[xz] \ni m$ and $[yz] \ni m$.

A very readable paper on the subject was written by Brian Bowditch [24].

6.24. Exercise. *Prove the implication (c)⇒(a) in the theorem.*

All the topics discussed in this lecture link Alexandrov geometry with the fundamental group. The theory of hyperbolic groups, a branch of geometric group theory, introduced by Mikhael Gromov [48], could be considered as a further step in this direction.

Lecture 7

Subsets

This lecture is nearly a copy of [8, Chapter 4]; here we give a partial answer to the following question: *Which subsets of Euclidean space, equipped with their induced length-metrics, are* CAT(0)*?*

A Motivating examples

Consider three subgraphs of different quadric surfaces:

$$\begin{aligned} A &= \left\{ (x, y, z) \in \mathbb{E}^3 \, : \, z \leqslant x^2 + y^2 \right\}, \\ B &= \left\{ (x, y, z) \in \mathbb{E}^3 \, : \, z \leqslant -x^2 - y^2 \right\}, \\ C &= \left\{ (x, y, z) \in \mathbb{E}^3 \, : \, z \leqslant x^2 - y^2 \right\}. \end{aligned}$$

7.1. Question. *Which of the sets A, B and C, if equipped with the induced length metric, are* CAT(0) *and why?*

The answers are given below, but it is instructive to think about these questions before reading further.

A. No, A is not CAT(0).

The boundary ∂A is the paraboloid described by $z = x^2 + y^2$; in particular it bounds an open convex set in $\mathbb{E}^3$ whose complement is A. The closest point projection of $A \to \partial A$ is short (Exercise 2.13). It follows that ∂A is a convex set in A equipped with its induced length metric.

Therefore if A is CAT(0), then so is ∂A. The latter is not true: ∂A is a smooth convex surface, and has strictly positive curvature by the Gauss formula.

B. Yes, B is CAT(0).

Evidently B is a convex closed set in $\mathbb{E}^3$. Therefore the length metric on B coincides with the Euclidean metric and CAT(0) comparison holds.

C. Yes, C is CAT(0), but the proof is not as easy as before.

Set $f_t(x,y) = x^2 - y^2 - 2{\cdot}(x-t)^2$. Consider the one-parameter family of sets

$$V_t = \left\{\,(x,y,z)\in\mathbb{E}^3\;:\;z\leqslant f_t(x,y)\,\right\}.$$

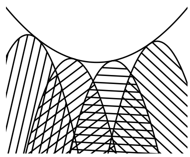

Each set V_t is a solid paraboloid tangent to ∂C along the parabola $y\mapsto(t,y,t^2-y^2)$. The set V_t is closed and convex for any t, and

$$C=\bigcup_t V_t.$$

Further note that the function $t\mapsto f_t(x,y)$ is concave for any fixed x,y. Therefore

❶ *if $a<b<c$, then $V_b\supset V_a\cap V_c$.*

Consider the finite union

$$C'=V_{t_1}\cup\dots\cup V_{t_n}.$$

The inclusion ❶ makes it possible to apply Reshetnyak gluing theorem 2.16 recursively and show that C' is CAT(0).

By approximation, the CAT(0) comparison holds for any 4 points in C; hence C is CAT(0). More precisely, choose $x_1,x_2,x_3,x_4\in C$ and 6 geodesics $[x_ix_j]_C$ between them. Choose $\varepsilon>0$, shift each $[x_ix_j]_C$ down by ε, and reconnect it to x_i and x_j by vertical ε-segments. Denote the obtained curve by $\gamma_{i,j}$; note that

$$\operatorname{length}\gamma_{i,j} = |x_i-x_j|_C + 2{\cdot}\varepsilon.$$

We may assume that C' contains all $\gamma_{i,j}$. It follows that

$$|x_i-x_j|_C\leqslant|x_i-x_j|_{C'}\leqslant|x_i-x_j|_C+2{\cdot}\varepsilon$$

Since and C' is CAT(0), $\varepsilon>0$ is arbitrary, so is C.

Remark. The set C is not convex, but it is *two-convex* as defined in the next section. As you will see, two-convexity is closely related to the inheritance of an upper curvature bound by a subset.

B Two-convexity

7.2. Definition. *We say that a subset $K \subset \mathbb{E}^m$ is* two-convex *if the following condition holds for any plane $W \subset \mathbb{E}^m$: If γ is a simple closed curve in $W \cap K$ that is null-homotopic in K, then it is null-homotopic in $W \cap K$, and in particular the disc in W bounded by γ lies in K.*

Note that two-convex sets do not have to be connected or simply connected. The following two propositions follow immediately from the definition.

7.3. Proposition. *Any subset in $\mathbb{E}^2$ is two-convex.*

7.4. Proposition. *The intersection of an arbitrary collection of two-convex sets in $\mathbb{E}^m$ is two-convex.*

7.5. Proposition. *The interior of any two-convex set in $\mathbb{E}^m$ is a two-convex set.*

Proof. Fix a two-convex set $K \subset \mathbb{E}^m$ and a 2-plane W; denote by $\operatorname{Int} K$ the interior of K. Let γ be a closed simple curve in $W \cap \operatorname{Int} K$ that is contractible in the interior of K.

Since K is two-convex, the plane disc D bounded by γ lies in K. The same holds for the translations of D by small vectors. Therefore D lies in $\operatorname{Int} K$; that is, $\operatorname{Int} K$ is two-convex. $\square$

7.6. Definition. *Given a subset $K \subset \mathbb{E}^m$, define its two-convex hull (briefly, $\operatorname{Conv}_2 K$) as the intersection of all two-convex subsets containing K.*

Note that by Proposition 7.4, the two-convex hull of any set is two-convex. Further, by 7.5, the two-convex hull of an open set is open.

The next proposition describes closed two-convex sets with smooth boundary.

7.7. Proposition. *Let $K \subset \mathbb{E}^m$ be a closed subset.*

Assume that the boundary of K is a smooth hypersurface S. Consider the unit normal vector field ν on S that points outside of K. Denote by $k_1 \leqslant \ldots \leqslant k_{m-1}$ the principal curvature functions of S with respect to ν (note that if K is convex, then $k_1 \geqslant 0$).

Then K is two-convex if and only if $k_2(p) \geqslant 0$ for any point $p \in S$. Moreover, if $k_2(p) < 0$ at some point p, then Definition 7.2 fails for some curve γ forming a triangle in an arbitrary small neighborhood of p.

The following proof was given by Mikhael Gromov [49, §$\tfrac12$], but we added a few details. The proof uses a straightforward modification of the Morse theory for manifolds with boundary; the paper of Sergei Vakhrameev [89] contains all the necessary lemmas.

Proof; only-if part. If $k_2(p) < 0$ for some $p \in S$, consider the plane W containing p and spanned by the first two principal directions at p. Choose a small triangle in W which surrounds p and move it slightly in the direction of $\nu(p)$. We get a triangle $[xyz]$ which is null-homotopic in K, but the solid triangle $\Delta = \operatorname{Conv}\{x, y, x\}$ bounded by $[xyz]$ does not lie in K completely. Therefore K is not two-convex. (See figure in the "only if" part of the smooth two-convexity theorem (7.10).)

If part. Recall that a smooth function $f\colon \mathbb{E}^m \to \mathbb{R}$ is called strongly convex if its Hessian is positive definite at each point.

Suppose $f\colon \mathbb{E}^m \to \mathbb{R}$ is a smooth strongly convex function such that the restriction $f|_S$ is a Morse function. Note that a generic smooth strongly convex function $f\colon \mathbb{E}^m \to \mathbb{R}$ has this property.

For a critical point p of $f|_S$, the outer normal vector $\nu(p)$ is parallel to the gradient $\nabla_p f$; we say that p is a positive critical point if $\nu(p)$ and $\nabla_p f$ point in the same direction, and negative otherwise. If f is generic, then we can assume that the sign is defined for all critical points; that is, $\nabla_p f \neq 0$ for any critical point p of $f|_S$.

Since $k_2 \geqslant 0$ and the function f is strongly convex, the negative critical points of $f|_S$ have index at most 1.

Given a real value s, set

$$K_s = \{\, x \in K \;:\; f(x) < s \,\}.$$

Assume $\varphi_0\colon \mathbb{D} \to K$ is a continuous map of the disc $\mathbb{D}$ such that $\varphi_0(\partial\mathbb{D}) \subset K_s$.

Note that by the Morse lemma, there is a homotopy $\varphi_t\colon \mathbb{D} \to K$ rel $\partial\mathbb{D}$ such that $\varphi_1(\mathbb{D}) \subset K_s$.

Indeed, we can construct a homotopy $\varphi_t\colon \mathbb{D} \to K$ that decreases the maximum of $f \circ \varphi$ on $\mathbb{D}$ until the maximum occurs at a critical point p of $f|_S$. This point cannot be negative; otherwise, its index would be at least 2. If this critical point is positive, then it is easy to decrease the maximum a little by pushing the disc from S into K in the direction of $-\nabla f_p$.

Consider a closed curve $\gamma\colon \mathbb{S}^1 \to K$ that is null-homotopic in K. Note that the distance function

$$f_0(x) = |\operatorname{Conv}\gamma - x|_{\mathbb{E}^m}$$

is convex. Therefore f_0 can be approximated by smooth strongly convex functions f in general position. From above, there is a disc in

K with boundary γ that lies arbitrarily close to $\operatorname{Conv}\gamma$. Since K is closed, the statement follows. $\square$

Note that the "if" part proves a somewhat stronger statement. Namely, any plane curve γ (not necessary simple) which is contractible in K is also contractible in the intersection of K with the plane of γ. The latter condition does not hold for the complement of two planes in $\mathbb{E}^4$, which is two-convex by Proposition 7.4; see also Exercise 7.18 below. The following proposition shows that there are no such examples in $\mathbb{E}^3$.

7.8. Proposition. *Let $\Omega \subset \mathbb{E}^3$ be an open two-convex subset. Then for any plane $W \subset \mathbb{E}^3$, any closed curve in $W \cap \Omega$ that is null-homotopic in Ω is also null-homotopic in $W \cap \Omega$.*

This statement is intuitively obvious, but the proof is not trivial; it use the following classical result. An alternative definition of two-convexity using homology instead of homotopy is mentioned in the last section. For this definition the proof is simpler.

7.9. Loop theorem. *Let M be a three-dimensional manifold with nonempty boundary ∂M. Assume $f\colon (\mathbb{D}, \partial\mathbb{D}) \to (M, \partial M)$ is a continuous map from the disc $\mathbb{D}$ such that the boundary curve $f|_{\partial\mathbb{D}}$ is not null-homotopic in ∂M. Then there is an embedding $h\colon (\mathbb{D}, \partial\mathbb{D}) \to (M, \partial M)$ with the same property.*

The theorem is due to Christos Papakyriakopoulos [a proof can be found in 53].

Proof of 7.8. Fix a closed plane curve γ in $W \cap \Omega$ that is null-homotopic in Ω. Suppose γ is not contractible in $W \cap \Omega$.

Let $\varphi\colon \mathbb{D} \to \Omega$ be a map of the disc with the boundary curve γ.

Since Ω is open we can first change φ slightly so that $\varphi(x) \notin W$ for $1 - \varepsilon < |x| < 1$ for some small $\varepsilon > 0$. By further changing φ slightly we can assume that it is transversal to W on $\operatorname{Int}\mathbb{D}$ and agrees with the previous map near $\partial\mathbb{D}$.

This means that $\varphi^{-1}(W) \cap \operatorname{Int}\mathbb{D}$ consists of finitely many simple closed curves which cut $\mathbb{D}$ into several components. Consider one of the "innermost" components c'; that is, c' is a boundary curve of a disc $\mathbb{D}' \subset \mathbb{D}$, $\varphi(c')$ is a closed curve in W and $\varphi(\mathbb{D}')$ completely lies in one of the two half-spaces with boundary W. Denote this half-space by H.

If $\varphi(c')$ is not contractible in $W \cap \Omega$, then applying the loop theorem to $M^3 = H \cap \Omega$ we conclude that there exists a simple closed curve $\gamma' \subset \Omega \cap W$ which is not contractible in $\Omega \cap W$ but is contractible in $\Omega \cap H$. This contradicts two-convexity of Ω.

Hence $\varphi(c')$ is contractible in $W \cap \Omega$. Therefore φ can be changed in a small neighborhood U of $\mathbb{D}'$ so that the new map $\hat{\varphi}$ maps U to one side of W. In particular, the set $\hat{\varphi}^{-1}(W)$ consists of the same curves as $\varphi^{-1}(W)$ with the exception of c'.

Repeating this process several times we reduce the problem to the case where $\varphi^{-1}(W) \cap \operatorname{Int}\mathbb{D} = \varnothing$. This means that $\varphi(\mathbb{D})$ lies entirely in one of the half-spaces bounded by W.

Again applying the loop theorem, we obtain a simple closed curve in $W \cap \Omega$ which is not contractible in $W \cap \Omega$ but is contractible in Ω. This again contradicts two-convexity of Ω. Hence γ is contractible in $W \cap \Omega$ as claimed. □

C Sets with smooth boundary

In this section, we characterize the subsets with smooth boundary in $\mathbb{E}^m$ that form CAT(0) spaces.

7.10. Smooth two-convexity theorem. *Let K be a closed, simply connected subset in $\mathbb{E}^m$ equipped with the induced length metric. Assume K is bounded by a smooth hypersurface. Then K is* CAT(0) *if and only if K is two-convex.*

This theorem is a baby case of a result of Stephanie Alexander, David Berg, and Richard Bishop [3], which is briefly discussed at the end of the lecture. The proof below is based on the argument in Section 7A.

Proof. Denote by S and by Ω the boundary and the interior of K respectively. Since K is connected and S is smooth, Ω is also connected.

Denote by $k_1(p) \leqslant \ldots \leqslant k_{m-1}(p)$ the principal curvatures of S at $p \in S$ with respect to the normal vector $\nu(p)$ pointing out of K. By Proposition 7.7, K is two-convex if and only if $k_2(p) \geqslant 0$ for any $p \in S$.

Only-if part. Assume K is not two-convex. Then by Proposition 7.7, there is a triangle $[xyz]$ in K which is null-homotopic in K, but the solid triangle $\Delta = \operatorname{Conv}\{x, y, z\}$ does not lie in K completely. Evidently the triangle $[xyz]$ is not thin in K. Hence K is not CAT(0).

If part. Since K is simply connected, by the globalization theorem (5.6) it suffices to show that any point $p \in K$ admits a CAT(0) neighborhood.

If $p \in \operatorname{Int} K$, then it admits a neighborhood isometric to a CAT(0) subset of $\mathbb{E}^m$. Fix $p \in S$. Assume that $k_2(p) > 0$. Fix a sufficiently small $\varepsilon > 0$ and set $K' = K \cap \overline{\mathrm{B}}[p, \varepsilon]$. Let us show that

❶ *K' is* CAT(0).

Consider the coordinate system with the origin at p and the principal directions and $\nu(p)$ as the coordinate directions. For small $\varepsilon > 0$, the set K' can be described as a subgraph

$$K' = \{ (x_1, \ldots, x_m) \in \overline{\mathrm{B}}[p, \varepsilon] \ : \ x_m \leqslant f(x_1, \ldots, x_{m-1} \} .$$

Fix $s \in [-\varepsilon, \varepsilon]$. Since ε is small and $k_2(p) > 0$, the restriction $f|_{x_1=s}$ is concave in the $(m-2)$-dimensional cube defined by the inequalities $|x_i| < 2{\cdot}\varepsilon$ for $2 \leqslant i \leqslant m-1$.

Fix a negative real value $\lambda < k_1(p)$. Given $s \in (-\varepsilon, \varepsilon)$, consider the set

$$V_s = \{ (x_1, \ldots, x_m) \in K' \ : \ x_m \leqslant f(x_1, \ldots, x_{m-1}) + \lambda{\cdot}(x_1 - s)^2 \} .$$

Note that the function

$$(x_1, \ldots, x_{m-1}) \mapsto f(x_1, \ldots, x_{m-1}) + \lambda{\cdot}(x_1 - s)^2$$

is concave near the origin. Since ε is small, we can assume that the V_s are convex subsets of $\mathbb{E}^m$.

Further note that

$$K' = \bigcup_{s \in [-\varepsilon, \varepsilon]} V_s.$$

Also, the same argument as in 7.1 shows that

❷ *If $a < b < c$, then $V_b \supset V_a \cap V_c$.*

Given an array of values $s^1 < \cdots < s^k$ in $[-\varepsilon, \varepsilon]$, set $V^i = V_{s^i}$ and consider the unions

$$W^i = V^1 \cup \cdots \cup V^i$$

equipped with the induced length metric.

Note that the array (s^n) can be chosen in such a way that W^k is arbitrarily close to K' in the sense of Hausdorff.

Arguing as in 7A, we get that the following claim implies ❶.

❸ *All W^i are* CAT(0).

This claim is proved by induction. Base: $W^1 = V^1$ is CAT(0) as a convex subset in $\mathbb{E}^m$.

Step: Assume that W^i is CAT(0). According to ❷,

$$V^{i+1} \cap W^i = V^{i+1} \cap V^i.$$

Moreover, this is a convex set in $\mathbb{E}^m$ and therefore it is a convex set in W^i and in V^{i+1}. By the Reshetnyak gluing theorem, W^{i+1} is CAT(0). Hence the claim follows. △

Note that we have proved the following:

❹ *K' is* CAT(0) *if K is strongly two-convex, that is, $k_2(p) > 0$ at any point $p \in S$.*

It remains to show that p admits a CAT(0) neighborhood in the case $k_2(p) = 0$.

Choose a coordinate system $(x_1, \dots, x_m)$ as above, so that the $(x_1, \dots, x_{m-1})$-coordinate hyperplane is the tangent subspace to S at p.

Fix $\varepsilon > 0$ so that a neighborhood of p in S is the graph

$$x_m = f(x_1, \dots, x_{m-1})$$

of a function f defined on the open ball B of radius ε centered at the origin in the $(x_1, \dots, x_{m-1})$-hyperplane. Fix a smooth positive strongly convex function $\varphi\colon B \to \mathbb{R}_+$ such that $\varphi(x) \to \infty$ as x approaches the boundary of B. Note that for $\delta > 0$, the subgraph K_δ defined by the inequality

$$x_m \leqslant f(x_1, \dots, x_{m-1}) - \delta \cdot \varphi(x_1, \dots, x_{m-1})$$

is strongly two-convex. By ❹, K_δ is CAT(0).

Finally as $\delta \to 0$, the closed ε-neighborhoods of p in K_δ converge to the closed ε-neighborhood of p in K. By 2.1, the ε-neighborhood of p is CAT(0). □

D Open plane sets

In this section, we consider inheritance of upper curvature bounds by subsets of the Euclidean plane.

7.11. Theorem. *Let Ω be an open simply connected subset of $\mathbb{E}^2$. Equip Ω with its induced length metric and denote its completion by K. Then K is* CAT(0).

The assumption that the set Ω is open is not critical; instead one can assume that the induced length metric takes finite values at all points of Ω. We sketch the proof given by Richard Bishop [21] and leave the details to be finished as an exercise. A generalization of this result is proved by Alexander Lytchak and Stefan Wenger [65, Proposition 12.1]; this paper also contains a far-reaching application.

Sketch of proof. It is sufficient to show that any triangle in K is thin, as defined in 2.7.

Note that K admits a length-preserving map to $\mathbb{E}^2$ that extends the embedding $\Omega \hookrightarrow \mathbb{E}^2$. Therefore each triangle $[xyz]$ in K can be mapped to the plane in a length-preserving way. Since Ω is simply connected, any open region, say Δ, that is surrounded by the image of $[xyz]$ lies completely in Ω.

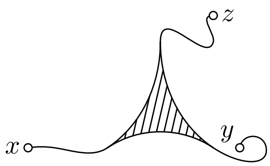

Note that in each triangle $[xyz]$ in K, the sides $[xy]$, $[yz]$ and $[zx]$ intersect each other along a geodesic starting at a common vertex, possibly a one-point geodesic. In other words, every triangle in K looks like the one in the diagram.

Indeed, assuming the contrary, there will be a *lune* in K bounded by two minimizing geodesics with common ends but no other common points. The image of this lune in the plane must have concave sides, since otherwise one could shorten the sides by pushing them into the interior. Evidently, there is no plane lune with concave sides, a contradiction.

Note that it is sufficient to consider only simple triangles $[xyz]$, that is, triangles whose sides $[xy]$, $[yz]$ and $[zx]$ intersect each other only at the common vertices. If this is not the case, chopping the overlapping part of sides reduces to the injective case (this is formally stated in Exercise 7.12).

Again, the open region, say Δ, bounded by the image of $[xyz]$ has concave sides in the plane, since otherwise one could shorten the sides by pushing them into Ω. It remains to solve Exercise 7.13. $\square$

7.12. Exercise. *Assume that $[pq]$ is a common part of the two sides $[px]$ and $[py]$ of the triangle $[pxy]$. Consider the triangle $[qxy]$ whose sides are formed by arcs of the sides of $[pxy]$. Show that if $[qxy]$ is thin, then so is $[pxy]$.*

7.13. Exercise. *Assume S is a closed plane region whose boundary is a plane triangle T with concave sides. Equip S with the induced length metric. Show that the triangle T is thin in S.*

Here is a spherical analog of Theorem 7.11, which can be proved along the same lines. It will be used in the next section.

7.14. Proposition. *Let Θ be an open connected subset of the unit sphere $\mathbb{S}^2$ that does not contain a closed hemisphere. Equip Θ with the induced length metric. Let $\tilde\Theta$ be a metric cover of Θ such that any closed curve in $\tilde\Theta$ shorter than $2\cdot\pi$ is contractible.*

Show that the completion of $\tilde\Theta$ is CAT(1).

7.15. Exercise. *Prove the following partial case of the proposition:*

Let K be closed subset of the unit sphere $\mathbb{S}^2$ that does not contain a closed hemisphere. Suppose K is simply connected and bounded by *a simple Lipschitz curve. Show that K with induced length metric is* CAT(1).

E Shefel's theorem

In this section, we will formulate our version of a theorem of Samuel Shefel (7.17) and prove a couple of its corollaries.

It seems that Shefel was very intrigued by the survival of metric properties under affine transformation. To describe an instance of such phenomena, note that two-convexity survives under affine transformations of a Euclidean space. Therefore, as a consequence of the smooth two-convexity theorem (7.10), the following holds.

7.16. Corollary. *Let K be closed connected subset of Euclidean space equipped with the induced length metric. Assume K is* CAT(0) *and bounded by a smooth hypersurface. Then any affine transformation of K is also* CAT(0).

By Corollary 7.19, an analogous statement holds for sets bounded by Lipschitz surfaces in the three-dimensional Euclidean space. In higher dimensions this is no longer true.

7.17. Two-convexity theorem. *Let Ω be a connected open set in $\mathbb{E}^3$. Equip Ω with the induced length metric and denote by $\tilde{K}$ the completion of the universal metric cover of Ω. Then $\tilde{K}$ is* CAT(0) *if and only if Ω is two-convex.*

The proof of this statement will be given in the following three sections. First we prove its polyhedral analog, then we prove some properties of two-convex hulls in three-dimensional Euclidean space and only then do we prove the general statement.

The following exercise shows that the analogous statement does not hold in higher dimensions.

7.18. Exercise. *Let Π_1, Π_2 be two planes in $\mathbb{E}^4$ intersecting at a single point. Let $\tilde{K}$ be the completion of the universal metric cover of $\mathbb{E}^4 \setminus (\Pi_1 \cup \Pi_2)$.*

Show that $\tilde{K}$ is CAT(0) *if and only if $\Pi_1 \perp \Pi_2$.*

Before coming to the proof of the two-convexity theorem, let us formulate a few corollaries. The following corollary is a generalization of the smooth two-convexity theorem (7.10) for three-dimensional Euclidean space.

7.19. Corollary. *Let K be a closed subset in $\mathbb{E}^3$ bounded by a Lipschitz hypersurface. Then K with the induced length metric is* CAT(0) *if and only if the interior of K is two-convex and simply connected.*

Proof. Set $\Omega = \operatorname{Int} K$. Since K is simply connected and bounded by a surface, Ω is also simply connected.

Apply the two-convexity theorem to Ω. Note that the completion of Ω equipped with the induced length metric is isometric to K with the induced length metric. Hence the result. □

Note that the Lipschitz condition is used just once to show that the completion of Ω is isometric to K with the induced length metric. This property holds for a wider class of hypersurfaces; for instance Alexander horned ball might have $\mathrm{CAT}(0)$ induced length metric.

Let U be an open set in $\mathbb{R}^2$. A continuous function $f\colon U \to \mathbb{R}$ is called s a d d l e if for any linear function $\ell\colon \mathbb{R}^2 \to \mathbb{R}$, the difference $f - \ell$ does not have local maxima or local minima in U. Equivalently, the open subgraph and epigraph of f

$$\left\{\, (x,y,z) \in \mathbb{E}^3 \;:\; z < f(x,y),\ (x,y) \in U \,\right\},$$
$$\left\{\, (x,y,z) \in \mathbb{E}^3 \;:\; z > f(x,y),\ (x,y) \in U \,\right\}$$

are two-convex.

7.20. Theorem. *Let $f\colon \mathbb{D} \to \mathbb{R}$ be a Lipschitz function which is saddle in the interior of the closed unit disc $\mathbb{D}$. Then the graph*

$$\Gamma = \left\{\, (x,y,z) \in \mathbb{E}^3 \;:\; z = f(x,y) \,\right\},$$

equipped with induced length metric is $\mathrm{CAT}(0)$.

Proof. Since the function f is Lipschitz, its graph Γ with the induced length metric is bi-Lipschitz equivalent to $\mathbb{D}$ with the Euclidean metric.

Consider the sequence of sets

$$K_n = \left\{\, (x,y,z) \in \mathbb{E}^3 \;:\; z \lesseqgtr f(x,y) \pm \tfrac{1}{n},\ (x,y) \in \mathbb{D} \,\right\}.$$

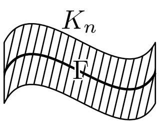

Note that each K_n is closed and simply connected. By definition K is also two-convex. Moreover the boundary of K_n is a Lipschitz surface.

Equip K_n with the induced length metric. By Corollary 7.19, K_n is $\mathrm{CAT}(0)$. It remains to note that $K_n \to \Gamma$ in the sense of Gromov–Hausdorff, and apply 2.1. □

F Polyhedral case

Now we are back to the proof of the two-convexity theorem (7.17).

Recall that a subset P of $\mathbb{E}^m$ is called a p o l y t o p e if it can be presented as a union of a finite number of simplices. Similarly, a s p h e r i c a l p o l y t o p e is a union of a finite number of simplices in $\mathbb{S}^m$.

Note that any polytope admits a finite triangulation. Therefore any polytope equipped with the induced length metric forms a Euclidean polyhedral space as defined in 6.4.

7.21. Lemma. *The two-convexity theorem (7.17) holds if the set Ω is the interior of a polytope.*

The statement might look obvious, but there is a hidden obstacle in the proof that is related to the following. Let P be a polytope and Ω its interior, both considered with the induced length metrics. Typically, the completion K of Ω is isometric to P — in this case, the lemma follows easily from 6.5.

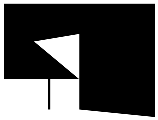

However in general we only have a locally distance-preserving map $K \to P$; it does not have to be onto and it may not be injective. An example can be guessed from the picture. Nevertheless, is easy to see that K is always a polyhedral space.

The proof uses the following two exercises.

7.22. Exercise. *Show that any closed path of length $< 2{\cdot}\pi$ in the units sphere $\mathbb{S}^2$ lies in an open hemisphere.*

7.23. Exercise. *Assume Ω is an open subset in $\mathbb{E}^3$ that is not two-convex. Show that there is a plane W such that the complement $W \setminus \Omega$ contains an isolated point and a small circle around this point in W is contractible in Ω.*

Proof of 7.21. The "only if" part can be proved in the same way as in the smooth two-convexity theorem (7.10) with additional use of Exercise 7.23.

If part. Assume that Ω is two-convex. Denote by $\tilde\Omega$ the universal metric cover of Ω. Let $\tilde K$ and K be the corresponding completions of $\tilde\Omega$ and Ω.

The main step is to show that $\tilde K$ is CAT(0).

Note that K is a polyhedral space and the covering $\tilde\Omega \to \Omega$ extends to a covering map $\tilde K \to K$ which might be branching at some vertices.[1]

Fix a point $\tilde p \in \tilde K \setminus \tilde\Omega$; denote by p the image of $\tilde p$ in K. Note that $\tilde K$ is a ramified cover of K and hence is locally contractible. Thus, any loop in $\tilde K$ is homotopic to a loop in $\tilde\Omega$ which is simply connected. Therefore $\tilde K$ is simply connected too.

[1]For example, if $K = \{\,(x,y,z) \in \mathbb{E}^3 \;:\; |z| \leqslant |x| + |y| \leqslant 1\,\}$ and p is the origin, then Σ_p, the space of directions at p, is not simply connected and $\tilde K \to K$ branches at p.

Thus, by the globalization theorem (5.6), it is sufficient to show that

❶ *a small neighborhood of $\tilde p$ in $\tilde K$ is* CAT(0).

Recall that $\Sigma_{\tilde p} = \Sigma_{\tilde p}\tilde K$ denotes the space of directions at $\tilde p$. Note that a small neighborhood of $\tilde p$ in $\tilde K$ is isometric to an open set in the cone over $\Sigma_{\tilde p}\tilde K$. By Exercise 6.2, ❶ follows once we can show that

❷ $\Sigma_{\tilde p}$ *is* CAT(1).

By rescaling, we can assume that every face of K which does not contain p lies at distance at least 2 from p. Denote by $\mathbb{S}^2$ the unit sphere centered at p, and set $\Theta = \mathbb{S}^2 \cap \Omega$. Note that $\Sigma_p K$ is isometric to the completion of Θ and $\Sigma_{\tilde p}\tilde K$ is the completion of the regular metric covering $\tilde\Theta$ of Θ induced by the universal metric cover $\tilde\Omega \to \Omega$.

By 7.14, it remains to show the following:

❸ *Any closed curve in $\tilde\Theta$ shorter than $2{\cdot}\pi$ is contractible.*

Fix a closed curve $\tilde\gamma$ of length $< 2{\cdot}\pi$ in $\tilde\Theta$. Its projection γ in $\Theta \subset \mathbb{S}^2$ has the same length. Therefore, by Exercise 7.22, γ lies in an open hemisphere. Then for a plane Π passing close to p, the central projection γ' of γ to Π is defined and lies in Ω. By construction of $\tilde\Theta$, the curve γ and therefore γ' are contractible in Ω. From two-convexity of Ω and Proposition 7.8, the curve γ' is contractible in $\Pi \cap \Omega$.

It follows that γ is contractible in Θ and therefore $\tilde\gamma$ is contractible in $\tilde\Theta$. □

G Two-convex hulls

The following proposition describes a construction which produces the two-convex hull $\operatorname{Conv}_2\Omega$ of an open set $\Omega \subset \mathbb{E}^3$. This construction is very close to the one given by Samuel Shefel [83].

7.24. Proposition. *Let $\Pi_1, \Pi_2 \dots$ be an everywhere dense sequence of planes in $\mathbb{E}^3$. Given an open set Ω, consider the recursively defined sequence of open sets $\Omega = \Omega_0 \subset \Omega_1 \subset \dots$ such that Ω_n is the union of Ω_{n-1} and all the bounded components of $\mathbb{E}^3 \setminus (\Pi_n \cup \Omega_{n-1})$. Then*

$$\operatorname{Conv}_2\Omega = \bigcup_n \Omega_n.$$

Proof. Set

❶ $$\Omega' = \bigcup_n \Omega_n.$$

Note that Ω' is a union of open sets; in particular, Ω' is open.

Let us show that

❷ $$\mathrm{Conv}_2\,\Omega \supset \Omega'.$$

Suppose we already know that $\mathrm{Conv}_2\,\Omega \supset \Omega_{n-1}$. Fix a bounded component $\mathfrak{C}$ of $\mathbb{E}^3\backslash(\Pi_n\cup\Omega_{n-1})$. It is sufficient to show that $\mathfrak{C}\subset\mathrm{Conv}_2\,\Omega$.

By 7.5, $\mathrm{Conv}_2\,\Omega$ is open. Therefore, if $\mathfrak{C}\not\subset\mathrm{Conv}_2\,\Omega$, then there is a point $p\in\mathfrak{C}\setminus\mathrm{Conv}_2\,\Omega$ lying at maximal distance from Π_n. Denote by W_p the plane containing p which is parallel to Π_n.

Note that p lies in a bounded component of $W_p\setminus\mathrm{Conv}_2\,\Omega$. In particular p can be surrounded by a simple closed curve γ in $W_p\cap\cap\,\mathrm{Conv}_2\,\Omega$. Since p lies at maximal distance from Π_n, the curve γ is null-homotopic in $\mathrm{Conv}_2\,\Omega$. Therefore $p\in\mathrm{Conv}_2\,\Omega$, a contradiction.

By induction, $\mathrm{Conv}_2\,\Omega\supset\Omega_n$ for each n. Therefore ❶ implies ❷.

It remains to show that Ω' is two-convex. Assume the contrary; that is, there is a plane Π and a simple closed curve $\gamma\colon\mathbb{S}^1\to\Pi\cap\Omega'$ which is null-homotopic in Ω', but not null-homotopic in $\Pi\cap\Omega'$.

By approximation we can assume that $\Pi=\Pi_n$ for a large n, and that γ lies in Ω_{n-1}. By the same argument as in the proof of Proposition 7.8 using the loop theorem, we can assume that there is an embedding $\varphi\colon\mathbb{D}\to\Omega'$ such that $\varphi|_{\partial\mathbb{D}}=\gamma$ and $\varphi(D)$ lies entirely in one of the half-spaces bounded by Π. By the n-step of the construction, the entire bounded domain U bounded by Π_n and $\varphi(D)$ is contained in Ω' and hence γ is contractible in $\Pi\cap\Omega'$, a contradiction. □

7.25. Key lemma. *The two-convex hull of the interior of a polytope in $\mathbb{E}^3$ is also the interior of a polytope.*

Proof. Fix a polytope P in $\mathbb{E}^3$. Set $\Omega=\mathrm{Int}\,P$. We may assume that Ω is dense in P (if not, redefine P as the closure of Ω). Denote by $F_1,\dots,F_m$ the facets of P. By subdividing F_i if necessary, we may assume that all F_i are convex polygons.

Set $\Omega'=\mathrm{Conv}_2\,\Omega$ and let P' be the closure of Ω'. Further, for each i, set $F_i'=F_i\setminus\Omega'$. In other words, F_i' is the subset of the facet F_i which remains on the boundary of P'.

From the construction of the two-convex hull (7.24) we have:

❸ *F_i' is a convex subset of F_i.*

Further, since Ω' is two-convex we obtain the following:

❹ *Each connected component of the complement $F_i\setminus F_i'$ is convex.*

Indeed, assume a connected component A of $F_i \setminus F_i'$ fails to be convex. Then there is a supporting line ℓ to F_i' touching F_i' at a single point in the interior of F_i. Then one could rotate the plane of F_i slightly around ℓ and move it parallelly to cut a "cap" from the complement of Ω. The latter means that Ω is not two-convex, a contradiction. $\triangle$

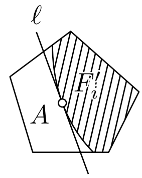

From ❸ and ❹, we conclude

❺ *F_i' is a convex polygon for each i.*

Consider the complement $\mathbb{E}^3 \setminus \Omega$ equipped with the length metric. By construction of the two-convex hull (7.24), the complement $L = \mathbb{E}^3 \setminus (\Omega' \cup P)$ is locally convex; that is, any point of L admits a convex neighborhood.

Summarizing: (1) Ω' is a two-convex open set, (2) the boundary $\partial\Omega'$ contains a finite number of polygons F_i' and the remaining part S of the boundary is locally concave. It remains to show that (1) and (2) imply that S and therefore $\partial\Omega'$ are piecewise linear.

7.26. Exercise. *Prove the last statement.* $\square$

H Proof of Shefel's theorem

Proof of 7.17. The "only if" part can be proved in the same way as in the smooth two-convexity theorem (7.10) with the additional use of Exercise 7.23.

If part. Suppose Ω is two-convex. We need to show that $\tilde{K}$ is CAT(0).

Fix a quadruple of points $x^1, x^2, x^3, x^4 \in \tilde{\Omega}$. Let us show that CAT(0) comparison holds for this quadruple.

Fix $\varepsilon > 0$. Choose six broken lines in $\tilde{\Omega}$ connecting all pairs of points x^1, x^2, x^3, x^4, where the length of each broken line is at most ε bigger than the distance between its ends in the length metric on $\tilde{\Omega}$. Denote by X the union of these broken lines. Choose a polytope P in Ω such that its interior $\operatorname{Int} P$ contains the projections of all six broken lines and discs which contract all the loops created by them (it is sufficient to take 3 discs).

Denote by Ω' the two-convex hull of the interior of P. According to the key lemma (7.25), Ω' is the interior of a polytope.

Equip Ω' with the induced length metric. Consider the universal metric cover $\tilde{\Omega}'$ of Ω'. (The covering $\tilde{\Omega}' \to \Omega'$ might be nontrivial — even if $\operatorname{Int} P$ is simply connected, its two-convex hull Ω' might not be simply connected.) Denote by $\tilde{K}'$ the completion of $\tilde{\Omega}'$.

By Lemma 7.21, $\tilde K'$ is CAT(0).

By construction of $\operatorname{Int} P$, the embedding $\operatorname{Int} P \hookrightarrow \Omega'$ admits a lift $\iota\colon X \hookrightarrow \tilde K'$. By construction, ι almost preserves the distances between the points x^1, x^2, x^3, x^4; namely

$$|\iota(x^i) - \iota(x^j)|_L \lessgtr |x^i - x^j|_{\operatorname{Int} P} \pm \varepsilon.$$

Since $\varepsilon > 0$ is arbitrary and CAT(0) comparison holds in $\tilde K'$, we get that CAT(0) comparison holds in Ω for x^1, x^2, x^3, x^4.

The statement follows since the quadruple $x^1, x^2, x^3, x^4 \in \tilde\Omega$ is arbitrary. $\square$

7.27. Exercise. *Assume $K \subset \mathbb{E}^m$ is a closed set bounded by a Lipschitz hypersurface. Equip K with the induced length metric. Show that if K is* CAT(0)*, then K is two-convex.*

The following exercise is analogous to Exercise 7.18. It provides a counterexample to the analog of Corollary 7.19 in higher dimensions.

7.28. Exercise. *Let $K = W \cap W'$, where*

$$W = \{\, (x, y, z, t) \in \mathbb{E}^4 \;:\; z \geqslant -x^2 \,\}$$

and $W' = \iota(W)$ for some motion $\iota\colon \mathbb{E}^4 \to \mathbb{E}^4$.

Show that K is always two-convex and one can choose ι so that K with the induced length metric is not CAT(0)*.*

I Remarks

Under the name $(n-2)$-convex sets, two-convex sets in $\mathbb{E}^n$ were introduced by Mikhael Gromov [49]. In addition to the inheritance of upper curvature bounds by two-convex sets discussed in this lecture, these sets appear as the maximal open sets with vanishing curvature in Riemannian manifolds with non-negative or non-positive sectional curvature [see Lemma 5.8 in 33, 17 and 72].

Two-convex sets could be defined using homology instead of homotopy, as in Gromov's formulation of the Leftschetz theorem [49, §$\tfrac12$]. Namely, we can say that K is two-convex if the following condition holds: if a one-dimensional cycle z has support in the intersection of K with a plane W and bounds in K, then it bounds in $K \cap W$.

The resulting definition is equivalent to the one used above. But unlike our definition it can be generalized to define k-convex sets in $\mathbb{E}^m$ for $k > 2$. With this homological definition one can also avoid the use of the loop theorem, whose proof is quite involved. Nevertheless, we chose the definition using homotopies since it is easier to visualize.

Both definitions work well for open sets; for general sets one should be able to give a similar definition using an appropriate homotopy/homology theory.

In [3], Stephanie Alexander, David Berg and Richard Bishop gave the exact upper bound on Alexandrov's curvature for the Riemannian manifolds with boundary. This theorem includes the smooth two-convexity theorem (7.10) as a partial case. Namely they show the following.

7.29. Theorem. *Let M be a Riemannian manifold with boundary ∂M. A direction tangent to the boundary will be called concave if there is a short geodesic in this direction which leaves the boundary and goes into the interior of M. A sectional direction (that is, a 2-plane) tangent to the boundary will be called concave if all the directions in it are concave.*

Denote by κ an upper bound of sectional curvatures of M and sectional curvatures of ∂M in the concave sectional directions. Then M is locally $\mathrm{CAT}(\kappa)$.

7.30. Corollary. *Let M be a Riemannian manifold with boundary ∂M. Assume that all the sectional curvatures of M and ∂M are bounded above by κ. Then M is locally* $\mathrm{CAT}(\kappa)$.

Theorem 7.20 is from Shefel's original paper [84]. It is related to Alexandrov's theorem about ruled surfaces [11].

Let D be an embedded closed disc in $\mathbb{E}^3$. We say that D is saddle if each connected component which any plane cuts from D contains a point on the boundary ∂D. If D is locally described by a Lipschitz embedding, then this condition is equivalent to saying that D is two-convex.

7.31. Shefel's conjecture. *Any saddle surface in $\mathbb{E}^3$ equipped with the length-metric is locally* $\mathrm{CAT}(0)$.

The conjecture is open even for the surfaces described by a bi-Lipschitz embedding of a disc. From another result of Samuel Shefel [84], it follows that a saddle surface satisfies the isoperimetric inequality $a \leqslant C\cdot\ell^2$ where a is the area of a disc bounded by a curve of length ℓ and $C = \frac{1}{3\cdot\pi}$. By a result of Alexander Lytchak and Stefan Wenger [65], Shefel's conjecture is equivalent to the isoperimetric inequality with the optimal constant $C = \frac{1}{4\cdot\pi}$. For more on the subject, see [64, 76] and the references therein.

Lecture 8

Barycenters

A Definition

Let us denote by $\triangle^k \subset \mathbb{R}^{k+1}$ the standard k-simplex; that is, $\boldsymbol{m} = (m_0, \dots, m_k) \in \triangle^k$ if $m_0 + \dots + m_k = 1$ and $m_i \geqslant 0$ for all i.

Consider a point array $\boldsymbol{p} = (p_0, \dots, p_k)$ in a Euclidean space $\mathbb{E}^n$. Recall that

$$z = m_0 \cdot p_0 + \dots + m_k \cdot p_k$$

is called barycenter of point array $\boldsymbol{p} = p_0, \dots, p_k$ with masses $\boldsymbol{m} = (m_0, \dots, m_k) \in \triangle^k$. Equivalently,

$$❶ \qquad z := \operatorname{MinPoint}(m_0 \cdot f_0 + \dots + m_k \cdot f_k),$$

where $f_i = \tfrac{1}{2} \cdot \operatorname{dist}^2_{p_i}$ for each i, and $\operatorname{MinPoint} f$ denotes a point of minimum of function f.

The map $\mathfrak{S} \colon \triangle^k \mapsto \mathbb{E}^n$ defined by $\mathfrak{S} \colon \boldsymbol{m} \mapsto z$ is called barycentric simplex of the array $\boldsymbol{p}$. If needed we may denote this map by $\mathfrak{S}_{\boldsymbol{p}}$ or, more generally, $\mathfrak{S}_{\boldsymbol{f}}$. The latter means that we define the map via ❶ for an array of functions $\boldsymbol{f} = (f_0, f_1, \dots, f_k)$. Formally speaking, the definition ❶ makes sense for any array of functions in a metric space; altho, the map might be undefined or nonuniquely defined.

Further, we will work with this definition in CAT(0) spaces instead of $\mathbb{E}^n$. It will be used to define and study dimension of CAT spaces. We will use that on a geodesic CAT(0) space, functions of the type $f = \tfrac{1}{2} \cdot \operatorname{dist}^2_p$ are 1-convex; see 2.10. Besides that, we will not use CAT(0) condition for a while.

B Barycentric simplex

8.1. Theorem. *Let $\mathcal{X}$ be a complete geodesic space and $\boldsymbol{f} = (f_0, \dots$ $\dots, f_k)\colon \mathcal{X} \to \mathbb{R}^{k+1}$ be an array of nonnegative 1-convex locally Lipschitz functions. Then the barycentric simplex $\mathfrak{S}_{\boldsymbol{f}}\colon \triangle^k \to \mathcal{X}$ is a uniquely defined Lipshitz map.*

In particular, we have that the barycentric simplex $\mathfrak{S}_{\boldsymbol{p}}$ any point array $\boldsymbol{p} = (p_0, \dots, p_k)$ in a complete geodesic $\mathrm{CAT}(0)$ *space is a uniquely defined Lipshitz map.*

8.2. Lemma. *Suppose $\mathcal{X}$ is a complete geodesic space and $f\colon \mathcal{X} \to \mathbb{R}$ is a locally Lipschitz,* 1*-convex function. Then* $\operatorname{MinPoint} f$ *is uniquely defined.*

Proof. Note that

❶ *if z is a midpoint of the geodesic $[xy]$, then*

$$s \leqslant f(z) \leqslant \tfrac{1}{2}\cdot f(x) + \tfrac{1}{2}\cdot f(y) - \tfrac{1}{8}\cdot|x-y|^2,$$

where s is the infimum of f.

Uniqueness. Assume that x and y are distinct minimum points of f. From ❶ we have

$$f(z) < f(x) = f(y)$$

— a contradiction.

Existence. Fix a point $p \in \mathcal{X}$, and let $L \in \mathbb{R}$ be a Lipschitz constant of f in a neighborhood of p.

Choose a sequence of points $p_n \in \mathcal{X}$ such that $f(p_n) \to s$. Applying ❶ for $x = p_n$, $y = p_m$, we see that p_n is a Cauchy sequence. Thus the sequence p_n converges to a minimum point of f. □

Proof of 8.1. Since each f_i is 1-convex, for any $\boldsymbol{x} = (x_0, x_1, \dots, x_k) \in$ $\in \triangle^k$ the convex combination

$$\left(\sum_i x_i\cdot f_i\right) : \mathcal{X} \to \mathbb{R}$$

is also 1-convex. Therefore, according to 8.2, the barycentric simplex $\mathfrak{S}_{\boldsymbol{f}}$ is uniquely defined on $\triangle^k$.

For $\boldsymbol{x}, \boldsymbol{y} \in \triangle^k$, let

$$\begin{aligned} f_{\boldsymbol{x}} &= \sum_i x_i\cdot f_i, & f_{\boldsymbol{y}} &= \sum_i y_i\cdot f_i, \\ p &= \mathfrak{S}_{\boldsymbol{f}}(\boldsymbol{x}), & q &= \mathfrak{S}_{\boldsymbol{f}}(\boldsymbol{y}), \end{aligned}$$

Choose a geodesic γ from p to q; suppose $s=|p-q|$ and so $\gamma(0)=p$ and $\gamma(s)=q$. Observe the following:

- ◇ The function $\varphi(t)=f_{\boldsymbol{x}}\circ\gamma(t)$ has a minimum at 0. Therefore $\varphi^+(0)\geqslant 0$.
- ◇ The function $\psi(t)=f_{\boldsymbol{y}}\circ\gamma(t)$ has a minimum at s. Therefore $\psi^-(s)\geqslant 0$.

From 1-convexity of $f_{\boldsymbol{y}}$, we have

$$\psi^+(0)+\psi^-(s)+s\leqslant 0.$$

Let L be a Lipschitz constant for all f_i in a neighborhood $\Omega\ni p$. Then

$$\psi^+(0)\leqslant\varphi^+(0)+L\cdot\|\boldsymbol{x}-\boldsymbol{y}\|_1,$$

where $\|\boldsymbol{x}-\boldsymbol{y}\|_1=\sum_{i=0}^k|x_i-y_i|$. It follows that given $\boldsymbol{x}\in\triangle^k$, there is a constant L such that

$$\begin{aligned}|\mathfrak{S}_{\boldsymbol{f}}(\boldsymbol{x})-\mathfrak{S}_{\boldsymbol{f}}(\boldsymbol{y})|&=s\leqslant\\&\leqslant L\cdot\|\boldsymbol{x}-\boldsymbol{y}\|_1\end{aligned}$$

for any $\boldsymbol{y}\in\triangle^k$. In particular, there is $\varepsilon>0$ such that if $\|\boldsymbol{x}-\boldsymbol{y}\|_1<\varepsilon$, $\|\boldsymbol{x}-\boldsymbol{z}\|_1<\varepsilon$, then $\mathfrak{S}_{\boldsymbol{f}}(\boldsymbol{y})$, $\mathfrak{S}_{\boldsymbol{f}}(\boldsymbol{z})\in\Omega$. Thus the same argument as above implies

$$|\mathfrak{S}_{\boldsymbol{f}}(\boldsymbol{y})-\mathfrak{S}_{\boldsymbol{f}}(\boldsymbol{z})|\leqslant L\cdot\|\boldsymbol{y}-\boldsymbol{z}\|_1$$

for any $\boldsymbol{y}$ and $\boldsymbol{z}$ sufficiently close to $\boldsymbol{x}$; that is, $\mathfrak{S}_{\boldsymbol{f}}$ is locally Lipschitz. Since $\triangle^k$ is compact, $\mathfrak{S}_{\boldsymbol{f}}$ is Lipschitz. □

8.3. Exercise. *Let G be a subgroup of the group of isometries of a proper geodesic* CAT(0) *space. Assume that*

(a) G is finite, or

(b) G is compact.

Show that the action of G has a fixed point.

C Convexity of up-set

8.4. Definition. *For two real arrays $\boldsymbol{v}$, $\boldsymbol{w}\in\mathbb{R}^{k+1}$, $\boldsymbol{v}=(v_0,\dots,v_k)$ and $\boldsymbol{w}=(w_0,\dots,w_k)$, we will write $\boldsymbol{v}\succcurlyeq\boldsymbol{w}$ if $v_i\geqslant w_i$ for each i.*

Given a subset $Q\subset\mathbb{R}^{k+1}$, denote by $\operatorname{Up}Q$ the smallest upper set containing Q; that is,

$$\operatorname{Up}Q=\left\{\,\boldsymbol{v}\in\mathbb{R}^{k+1}\ :\ \exists\,\boldsymbol{w}\in Q\text{ such that }\boldsymbol{v}\succcurlyeq\boldsymbol{w}\,\right\},$$

8.5. Proposition. *Let $\mathcal{X}$ be a complete geodesic space and $\boldsymbol{f}=(f_0,\dots,f_k)\colon\mathcal{X}\to\mathbb{R}^{k+1}$ be an array of nonnegative 1-convex locally Lipschitz functions. Consider the set $W=\operatorname{Up}[\boldsymbol{f}(\mathcal{X})]\subset\mathbb{R}^{k+1}$. Then* *(a)* *The set* W *is convex.*

(b) $\boldsymbol{f}[\mathfrak{S}_{\boldsymbol{f}}(\triangle^k)] \subset \partial W$*. Moreover,* $\boldsymbol{f}[\mathfrak{S}_{\boldsymbol{f}}(\triangle^k) \setminus \mathfrak{S}_{\boldsymbol{f}}(\partial\triangle^k)]$ *is an open set in* ∂W*.*

(c) $W = \operatorname{Up}(\boldsymbol{f}[\mathfrak{S}_{\boldsymbol{f}}(\triangle^k)])$*; in other words,* $\operatorname{Up}(\boldsymbol{f}[\mathfrak{S}_{\boldsymbol{f}}(\triangle^k)]) \supset \boldsymbol{f}(\mathcal{X})$*.*

Note that since $\triangle^k$ is compact, we also get that W is closed.

Proof. Let $V = \boldsymbol{f}(\mathcal{X}) \subset \mathbb{R}^{k+1}$; so $W = \operatorname{Up} V$. Denote by $\bar{V}$ the closure of V.

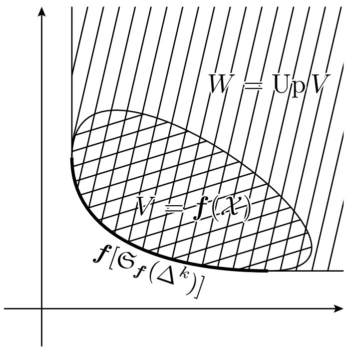

(a). Convexity of all f_i implies that for any two points $p, q \in \mathcal{X}$ and $t \in \in [0,1]$ we have

$$(1-t)\cdot\boldsymbol{f}(p) + t\cdot\boldsymbol{f}(q) \succcurlyeq \boldsymbol{f} \circ \gamma(t),$$

where γ denotes a geodesic path from p to q. Therefore, W is convex.

(b)+(c). Choose $p \in \mathfrak{S}_{\boldsymbol{f}}(\triangle^k)$. Note that if $\boldsymbol{f}(p) \succcurlyeq \boldsymbol{w}$ for some $\boldsymbol{w} \in W$, then $\boldsymbol{f}(p) = \boldsymbol{w}$. It follows that $\boldsymbol{f}(p) \in \partial W$; therefore $\boldsymbol{f}[\mathfrak{S}_{\boldsymbol{f}}(\triangle^k)]$ lies in a convex hypersurface ∂W.

Choose $\boldsymbol{w} \in W$. Observe that $\boldsymbol{w} \succcurlyeq \boldsymbol{v}$ for some $\boldsymbol{v} \in \bar{V} \cap \partial W$. Note that W is supported at $\boldsymbol{v}$ by a hyperplane

$$\Pi = \left\{ (x_1, \dots, x_k) \in \mathbb{R}^k \ : \ m_0\cdot x_0 + \dots + m_k\cdot x_k = \text{const} \right\}$$

for some $\boldsymbol{m} = (m_0, \dots, m_k) \in \triangle^k$. Let $p = \mathfrak{S}_{\boldsymbol{f}}(\boldsymbol{m})$. By 8.2, $\boldsymbol{f}(p) = \boldsymbol{v}$; in particular $\boldsymbol{v} \in V$.

Note that $p \in \mathfrak{S}_{\boldsymbol{f}}(\triangle^k) \setminus \mathfrak{S}_{\boldsymbol{f}}(\partial\triangle^k)$ if and only if $\boldsymbol{f}(p)$ is supported by a plane as above for some $\boldsymbol{m} \in \triangle^k$, but it is not supported by a plane for some $\boldsymbol{m} \in \partial\triangle^k$. This condition is open, therefore $\mathfrak{S}_{\boldsymbol{f}}(\triangle^k) \setminus \setminus \mathfrak{S}_{\boldsymbol{f}}(\partial\triangle^k)$ is an open set. □

D Nondegenerate simplex

Given an array $\boldsymbol{f} = (f_0, \dots, f_k)$, we denote by $\boldsymbol{f}^{-i}$ the subarray of $\boldsymbol{f}$ with f_i removed; that is,

$$\boldsymbol{f}^{-i} := (f_0, \dots, f_{i-1}, f_{i+1}, \dots, f_k).$$

It should be clear from the definition that $\mathfrak{S}_{\boldsymbol{f}^{-i}}$ coincides with the restriction of $\mathfrak{S}_{\boldsymbol{f}}$ to the corresponding facet of $\triangle^k$.

If $\operatorname{Im}\mathfrak{S}_{\boldsymbol{f}}$ is not covered by $\operatorname{Im}\mathfrak{S}_{\boldsymbol{f}^{-i}}$ for all i, then we say that $\mathfrak{S}_{\boldsymbol{f}}$ is nondegenerate. In other words, $\mathfrak{S}_{\boldsymbol{f}}$ is nondegenerate if

$$\mathfrak{S}_{\boldsymbol{f}}(\triangle^k)\setminus\mathfrak{S}_{\boldsymbol{f}}(\partial\triangle^k)\neq\varnothing.$$

8.6. Exercise. *Let $\mathcal{U}$ be a complete geodesic* CAT(0) *space.*

Show that the image 1-dimensional barycentric simplex for a pair of points $p_0, p_1\in\mathcal{U}$ is the geodesic $[p_0p_1]$.

Construct a CAT(0) *space with a three-point array (p_0,p_1,p_2) such that its barycentric simplex is nondegenerate and noninjective.*

8.7. Exercise. *Let $\boldsymbol{p}=(p_0,\dots,p_k)$ be a point array in a complete length* CAT(0) *space $\mathcal{U}$, and $B_i=\overline{\mathrm{B}}[p_i,r_i]$ for some array of positive reals $(r_0,r_1,\dots,r_k)$.*

(a) Suppose $\bigcap_i B_i\neq\varnothing$. Show that

$$\operatorname{Im}\mathfrak{S}_{\boldsymbol{p}}\subset\bigcup_i B_i.$$

(b) Suppose $\bigcap_i B_i=\varnothing$, but $\bigcap_{i\neq j} B_i\neq\varnothing$ for any j. Show that $\mathfrak{S}_{\boldsymbol{p}}$ is nondegenerate.

(c) Suppose $\mathfrak{S}_{\boldsymbol{p}}$ is nondegenerate. Show that the condition in (b) hold for some array of positive reals $(r_0,\dots,r_k)$.

E bi-Hölder embedding

8.8. Theorem. *Let $\mathcal{X}$ be a complete geodesic space and $\boldsymbol{f}=(f_0,\dots \dots,f_k)\colon\mathcal{X}\to\mathbb{R}^{k+1}$ be an array of 1-convex locally Lipschitz functions. Then the set*

$$Z=\mathfrak{S}_{\boldsymbol{f}}(\triangle^k)\setminus\mathfrak{S}_{\boldsymbol{f}}(\partial\triangle^k)$$

is $C^{\frac12}$-bi-Hölder to an open domain in $\mathbb{R}^k$.

Proof. Let $\operatorname{proj}\colon\mathbb{R}^{k+1}\to\Pi$ be orthogonal projection to the hyperplane $x_0+\dots+x_k=0$. Let us show that the restriction $\operatorname{proj}\circ\boldsymbol{f}|_Z$ is a bi-Hölder embedding.

The map $\operatorname{proj}\circ\boldsymbol{f}$ is Lipschitz; it remains to construct its right inverse and show that it is $C^{\frac12}$-continuous.

Given $\boldsymbol{v}=(v_0,v_1,\dots,v_k)\in\Pi$, consider the function $h_{\boldsymbol{v}}\colon\mathcal{X}\to\mathbb{R}$ defined by

$$h_{\boldsymbol{v}}(p)=\max_i\{f_i(p)-v_i\}.$$

Note that $h_{\boldsymbol{v}}$ is 1-convex. Let

$$\Phi(\boldsymbol{v}) := \operatorname{MinPoint} h_{\boldsymbol{v}}.$$

According to Lemma 8.2, $\Phi(\boldsymbol{v})$ is uniquely defined.

If $\boldsymbol{v} = \operatorname{proj} \boldsymbol{f}(p)$, then

$$f_i \circ \Phi(\boldsymbol{v}) \leqslant f_i(p)$$

for any i. In particular, if $p \in \mathfrak{S}_{\boldsymbol{f}}(\triangle^k)$, then $p = \Phi(\boldsymbol{v})$. That is, Φ is a right inverse of the restriction $\boldsymbol{f}|_{\mathfrak{S}_{\boldsymbol{f}}(\triangle^k)}$.

Given $\boldsymbol{v}, \boldsymbol{w} \in \mathbb{R}^{k+1}$, set $p = \Phi(\boldsymbol{v})$ and $q = \Phi(\boldsymbol{w})$. Since $h_{\boldsymbol{v}}$ and $h_{\boldsymbol{w}}$ are 1-convex, we have

$$h_{\boldsymbol{v}}(q) \geqslant h_{\boldsymbol{v}}(p) + \tfrac{1}{2}\cdot|p-q|^2, \qquad h_{\boldsymbol{w}}(p) \geqslant h_{\boldsymbol{w}}(q) + \tfrac{1}{2}\cdot|p-q|^2.$$

Therefore,

$$\begin{aligned}|p-q|^2 &\leqslant 2\cdot \sup_{x\in\mathcal{X}}\{|h_{\boldsymbol{v}}(x) - h_{\boldsymbol{w}}(x)|\} \leqslant \\ &\leqslant 2\cdot \max_i\{|v_i - w_i|\}.\end{aligned}$$

In particular, Φ is $C^{\frac{1}{2}}$-continuous.

Finally, by 8.5b, $\boldsymbol{f}(Z)$ is a k-dimensional manifold, hence the result. □

F Topological dimension

Let $\mathcal{X}$ be a metric space and $\{V_\beta\}_{\beta\in\mathcal{B}}$ be an open cover of $\mathcal{X}$. Let us recall two notions in general topology:

- ◇ The order of $\{V_\beta\}$ is the supremum of all integers n such that there is a collection of $n+1$ elements of $\{V_\beta\}$ with nonempty intersection.
- ◇ An open cover $\{W_\alpha\}_{\alpha\in\mathcal{A}}$ of $\mathcal{X}$ is called a refinement of $\{V_\beta\}_{\beta\in\mathcal{B}}$ if for any $\alpha \in \mathcal{A}$ there is $\beta \in \mathcal{B}$ such that $W_\alpha \subset V_\beta$.

8.9. Definition. *Let $\mathcal{X}$ be a metric space. The topological dimension of $\mathcal{X}$ is defined to be the minimum of nonnegative integers n such that for any open cover of $\mathcal{X}$ there is a finite open refinement with order n.*

If no such n exists, the topological dimension of $\mathcal{X}$ is infinite.

The topological dimension of $\mathcal{X}$ will be denoted by $\operatorname{TopDim}\mathcal{X}$.

The invariants satisfying the following two statements 8.10 and 8.11 are commonly called "dimension"; for that reason, we call these statements axioms.

8.10. Normalization axiom. *For any $m\in\mathbb{Z}_{\geqslant0}$,*

$$\operatorname{TopDim}\mathbb{E}^m=m.$$

8.11. Cover axiom. *If $\{A_n\}_{n=1}^\infty$ is a countable closed cover of $\mathcal{X}$, then*

$$\operatorname{TopDim}\mathcal{X}=\sup_n\{\operatorname{TopDim}A_n\}.$$

On product spaces. The following inequality holds for arbitrary metric spaces

$$\operatorname{TopDim}(\mathcal{X}\times\mathcal{Y})\leqslant\operatorname{TopDim}\mathcal{X}+\operatorname{TopDim}\mathcal{Y}.$$

It is strict for a pair of Pontryagin surfaces [77].

8.12. Definition. *Let $\mathcal{X}$ be a metric space and $F\colon\mathcal{X}\to\mathbb{R}^m$ be a continuous map. A point $\boldsymbol{z}\in\operatorname{Im}F$ is called a stable value of F if there is $\varepsilon>0$ such that $\boldsymbol{z}\in\operatorname{Im}F'$ for any ε-close to F continuous map $F'\colon\mathcal{X}\to\mathbb{R}^m$, that is, $|F'(x)-F(x)|<\varepsilon$ for all $x\in\mathcal{X}$.*

The next theorem follows from [55, theorems VI 1&2]. (This theorem also holds for non-separable metric spaces [70], [45, 3.2.10]).

8.13. Stable value theorem. *Let $\mathcal{X}$ be a separable metric space. Then $\operatorname{TopDim}\mathcal{X}\geqslant m$ if and only if there is a map $F\colon\mathcal{X}\to\mathbb{R}^m$ with a stable value.*

G Dimension theorem

8.14. Theorem. *For any proper geodesic $\mathrm{CAT}(0)$ space $\mathcal{U}$, the following statements are equivalent:*

(a)

$$\operatorname{TopDim}\mathcal{U}\geqslant m.$$

(b) For some $z\in\mathcal{U}$ there is an array of $m+1$ balls $B_i=\mathrm{B}(a_i,r_i)$ such that

$$\bigcap_i B_i=\varnothing\quad\text{and}\quad\bigcap_{i\neq j}B_i\neq\varnothing\quad\text{for each } j.$$

(c) There is a $C^{\frac12}$-embedding of an open set in $\mathbb{R}^m$ to $\mathcal{U}$; that is, Φ is bi-Hölder with exponent $\frac12$.

8.15. Lemma. *Let $\mathcal{U}$ be a proper geodesic* CAT(0) *space and $\rho\colon \mathcal{U} \to \mathbb{R}$ be a continuous positive function. Then there is a locally finite countable simplicial complex $\mathcal{N}$, a locally Lipschitz map $\Phi\colon \mathcal{U} \to \mathcal{N}$, and a Lipschitz map $\Psi\colon \mathcal{N} \to \mathcal{U}$ such that:*

(a) The displacement of the composition $\Psi \circ \Phi\colon \mathcal{U} \to \mathcal{U}$ is bounded by ρ; that is,

$$|x - \Psi \circ \Phi(x)| < \rho(x)$$

for any $x \in \mathcal{U}$.

(b) If $\operatorname{TopDim} \mathcal{U} \leqslant m$*, then the Ψ-image of $\mathcal{N}$ coincides with the image of its m-skeleton.*

Proof. Choose a locally finite countable covering $\{\,\Omega_\alpha : \alpha \in \mathcal{A}\,\}$ of $\mathcal{U}$ such that $\Omega_\alpha \subset \mathrm{B}(x, \tfrac{1}{3}\cdot\rho(x))$ for any $x \in \Omega_\alpha$.

Denote by $\mathcal{N}$ the nerve of the covering $\{\Omega_\alpha\}$; that is, $\mathcal{N}$ is an abstract simplicial complex with vertex set $\mathcal{A}$, such that a finite subset $\{\alpha_0, \dots, \alpha_n\} \subset \mathcal{A}$ forms a simplex if and only if

$$\Omega_{\alpha_0} \cap \dots \cap \Omega_{\alpha_n} \neq \varnothing.$$

Choose a Lipschitz partition of unity $\varphi_\alpha\colon \mathcal{U} \to [0,1]$ subordinate to $\{\Omega_\alpha\}$. Consider the map $\Phi\colon \mathcal{U} \to \mathcal{N}$ such that the barycentric coordinate of $\Phi(p)$ is $\varphi_\alpha(p)$. Note that Φ is locally Lipschitz. Clearly, the Φ-preimage of any open simplex in $\mathcal{N}$ lies in Ω_α for some $\alpha \in \mathcal{A}$.

For each $\alpha \in \mathcal{A}$, choose $x_\alpha \in \Omega_\alpha$. Let us extend the map $\alpha \mapsto x_\alpha$ to a map $\Psi\colon \mathcal{N} \to \mathcal{U}$ that is barycentric on each simplex. According to 8.1, this extension exists, and Ψ is locally Lipschitz.

(a). Fix $x \in \mathcal{U}$. Denote by $\triangle$ the minimal simplex that contains $\Phi(x)$, and let $\alpha_0, \alpha_1, \dots, \alpha_n$ be the vertexes of $\triangle$. Note that α is a vertex of $\triangle$ if and only if $\varphi_\alpha(x) > 0$. Thus

$$|x - x_{\alpha_i}| < \tfrac{1}{3}\cdot\rho(x)$$

for any i. Therefore

$$\operatorname{diam} \Psi(\triangle) \leqslant \max_{i,j}\{|x_{\alpha_i} - x_{\alpha_j}|\} < \tfrac{2}{3}\cdot\rho(x).$$

In particular,

$$|x - \Psi \circ \Phi(x)| \leqslant |x - x_{\alpha_0}| + \operatorname{diam} \Psi(\triangle) < \rho(x).$$

(b). Assume the contrary; that is, $\Psi(\mathcal{N})$ is not included in the Ψ-image of the m-skeleton of $\mathcal{N}$. Then for some $k > m$, there is a k-simplex $\triangle^k$

in $\mathcal{N}$ such that the barycentric simplex $\sigma = \Psi|_{\triangle^k}$ is nondegenerate; that is,

$$W = \Psi(\triangle^k) \setminus \Psi(\partial\triangle^k) \neq \varnothing.$$

By 8.8, $\operatorname{TopDim}\mathcal{U} \geqslant k$ — a contradiction. □

Proof of 8.14; (b)⇒(c)⇒(a). The implication *(b)⇒(c)* follows from Lemma 8.7 and Theorem 8.8, and *(c)⇒(a)* is trivial.

(a)⇒(b). According to 8.13, there is a continuous map $F\colon \mathcal{U} \to \mathbb{R}^m$ with a stable value.

Fix $\varepsilon > 0$. Since F is continuous, there is a continuous positive function ρ defined on $\mathcal{U}$ such that

$$|x - y| < \rho(x) \quad \Rightarrow \quad |F(x) - F(y)| < \tfrac{1}{3}{\cdot}\varepsilon.$$

Apply 8.15 to ρ. For the resulting simplicial complex $\mathcal{N}$ and the maps $\Phi\colon \mathcal{U} \to \mathcal{N}$, $\Psi\colon \mathcal{N} \to \mathcal{U}$, we have

$$|F \circ \Psi \circ \Phi(x) - F(x)| < \tfrac{1}{3}{\cdot}\varepsilon$$

for any $x \in \mathcal{U}$.

Arguing by contradiction, assume $\operatorname{TopDim}\mathcal{U} < m$. By 8.15*b*, the image $F_\varepsilon \circ \Psi \circ \Phi(K)$ lies in the F_ε-image of the $(m-1)$-skeleton of $\mathcal{N}$; In particular, it can be covered by a countable collection of Lipschitz images of $(m-1)$-simplexes. Hence $\mathbf{0} \in \mathbb{R}^m$ is not a stable value of $F_\varepsilon \circ \Psi \circ \Phi$. Since $\varepsilon > 0$ is arbitrary, we get the result. □

The following exercise is a generalization of Helly's theorem; for closely related statements see [59, Prop. 5.3] and [56].

8.16. Exercise. *Let $K_1, \dots, K_n$ be closed convex subsets in a proper length* CAT(0) *space $\mathcal{U}$. Suppose that* $\operatorname{TopDim}\mathcal{U} = m$ *and any $m+1$ subsets from $\{K_1, \dots, K_n\}$ have a common point. Show that all subsets $K_1, \dots, K_n$ have a common point.*

H Hausdorff dimension

8.17. Definition. *Let $\mathcal{X}$ be a metric space. Its Hausdorff dimension is defined as*

$$\operatorname{HausDim}\mathcal{X} = \sup\{\, \alpha \in \mathbb{R} \;:\; \operatorname{HausMes}_\alpha(\mathcal{X}) > 0 \,\},$$

where $\operatorname{HausMes}_\alpha$ *denotes the α-dimensional Hausdorff measure.*

The following theorem follows from [55, theorems V 8 and VII 2].

8.18. Szpilrajn's theorem. *Let $\mathcal{X}$ be a separable metric space. Assume* $\operatorname{TopDim}\mathcal{X} \geqslant m$. *Then* $\operatorname{HausMes}_m \mathcal{X} > 0$.

In particular, $\operatorname{TopDim}\mathcal{X} \leqslant \operatorname{HausDim}\mathcal{X}$.

Except for Szpilrajn's theorem, there are no other relations between topological and Husdorff dimension of separable spaces. Moreover, the following exercise implies that the same holds for compact geodesic CAT(0) spaces of topological dimension at least 1.

8.19. Exercise. *Construct a metric on the binary tree such that it has compact completion of arbitrary Huasdorff dimension $\alpha \geqslant 1$.*

Concclude that for any integer $m \geqslant 1$ and real $\alpha \geqslant m$ there is a compact CAT(0) *space with topological dimension m and Hausdorff dimension α.*

I Remarks

The barycenters in CAT(κ) spaces were introduced by Bruce Kleiner [59]. He also proved the dimension theorem; an improvement was made by Alexander Lytchak [66].

It is not known if the dimension theorem holds for arbitrary complete geodesic CAT(κ) spaces. It was conjectured by Bruce Kleiner [59], see also [50, p. 133]. For separable spaces, the answer is "yes", and it follows from Kleiner's argument [9, Corollary 14.13].

One may wonder if bi-Hölder condition 8.14*c* can be improved to bi-Lipschitz; it seem to be unknown even for compact spaces. However if a compact geodesic CAT(0) space $\mathcal{U}$ has finite topological dimension m, then a slight modification of Kleiner's technique can be used to show that there is a bi-Lipschitz embedding of an m-cube into $\mathcal{U}$ [9, Theorem 14.15]. In particular, there is a bi-Lipschitz embedding of an n-cube for any $n \leqslant m$. If $\operatorname{TopDim}\mathcal{U} = \infty$, then we expect existence of a bi-Lipschitz embedding of an n-cube for any integer $n \geqslant 1$. The statement is trivial for $n = 1$; in this case any geodesic gives an isometric embedding. For $n = 2$, one can get it from the the fact that minimal (or metric minimizing) surfaces in $\mathcal{U}$ are CAT(0) (any such surface is locally bi-Lipschitz to the Euclidean plane). For $n \geqslant 3$ the question remains open.

Semisolutions

0.1. Let $\mathcal{X}$ be a 4-point metric space.

Fix a tetrahedron $\triangle$ in $\mathbb{R}^3$. The vertices of $\triangle$, say x_0, x_1, x_2, x_3, can be identified with the points of $\mathcal{X}$.

Note that there is a unique quadratic form W on $\mathbb{R}^3$ such that

$$W(x_i - x_j) = |x_i - x_j|^2_{\mathcal{X}}$$

for all i and j.

By the triangle inequality, $W(v) \geqslant 0$ for any vector v parallel to one of the faces of $\triangle$.

Note that $\mathcal{X}$ is isometric to a 4-point subset in the Euclidean space if and only if $W(v) \geqslant 0$ for any vector v in $\mathbb{R}^3$.

Therefore, if $\mathcal{X}$ is not of type $\mathcal{E}_4$, then $W(v) < 0$ for some vector v. From above, the vector v must be transversal to each of the 4 faces of $\triangle$. Therefore if we project $\triangle$ along v to a plane transversal to v we see one of the two pictures on the right.

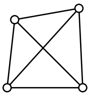

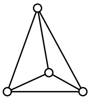

Note that the set of vectors v such that $W(v) < 0$ has two connected components; the opposite vectors v and $-v$ lie in the different components. If one moves v continuously, keeping $W(v) < 0$, then the corresponding projection moves continuously and the projections of the 4 faces cannot degenerate. It follows that the combinatorics of the picture do not depend on the choice of v. Hence $\mathcal{M}_4 \setminus \mathcal{E}_4$ is not connected.

It remains to show that if the combinatorics of the pictures for two spaces is the same, then one can continuously deform one space into the other. This can be easily done by deforming W and applying a permutation of x_0, x_1, x_2, x_3 if necessary.

Comment. This solution is inspired by [75].

0.2. The simplest proof we know requires the construction of tangent cones.

1.3. Consider the unit ball (B,ρ_0) in the space c_0 of all sequences converging to zero equipped with the sup-norm.

Consider another metric ρ_1 which is different from ρ_0 by the conformal factor

$$\varphi(\boldsymbol{x}) = 2 + \tfrac{1}{2}{\cdot}x_1 + \tfrac{1}{4}{\cdot}x_2 + \tfrac{1}{8}{\cdot}x_3 + \dots,$$

where $\boldsymbol{x} = (x_1, x_2 \dots) \in B$. That is, if $\boldsymbol{x}(t)$, $t \in [0,\ell]$, is a curve parametrized by ρ_0-length then its ρ_1-length is

$$\operatorname{length}_{\rho_1} \boldsymbol{x} = \int_0^\ell \varphi\circ\boldsymbol{x}.$$

Note that the metric ρ_1 is bi-Lipschitz to ρ_0.

Assume $\boldsymbol{x}(t)$ and $\boldsymbol{x}'(t)$ are two curves parametrized by ρ_0-length that differ only in the m-th coordinate, denoted by $x_m(t)$ and $x'_m(t)$ respectively. Note that if $x'_m(t) \leqslant x_m(t)$ for any t and the function $x'_m(t)$ is locally 1-Lipschitz at all t such that $x'_m(t) < x_m(t)$, then

$$\operatorname{length}_{\rho_1} \boldsymbol{x}' \leqslant \operatorname{length}_{\rho_1} \boldsymbol{x}.$$

Moreover this inequality is strict if $x'_m(t) < x_m(t)$ for some t.

Fix a curve $\boldsymbol{x}(t)$, $t \in [0,\ell]$, parametrized by ρ_0-length. We can choose m large, so that $x_m(t)$ is sufficiently close to 0 for any t. In particular, for some values t, we have $y_m(t) < x_m(t)$, where

$$y_m(t) = (1-\tfrac{t}{\ell}){\cdot}x_m(0) + \tfrac{t}{\ell}{\cdot}x_m(\ell) - \tfrac{1}{100}\cdot\min\{t, \ell - t\}.$$

Consider the curve $\boldsymbol{x}'(t)$ as above with

$$x'_m(t) = \min\{x_m(t), y_m(t)\}.$$

Note that $\boldsymbol{x}'(t)$ and $\boldsymbol{x}(t)$ have the same end points, and by the above

$$\operatorname{length}_{\rho_1} \boldsymbol{x}' < \operatorname{length}_{\rho_1} \boldsymbol{x}.$$

That is, for any curve $\boldsymbol{x}(t)$ in (B,ρ_1), we can find a shorter curve $\boldsymbol{x}'(t)$ with the same end points. In particular, (B,ρ_1) has no geodesics.

Comment. This example is due to Fedor Nazarov [71].

1.7. Consider the following subset of $\mathbb{R}^2$ equipped with the induced length metric

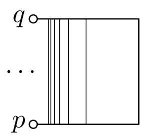

$$\mathcal{X} = \big((0,1]\times\{0,1\}\big) \cup \big(\{1, \tfrac{1}{2}, \tfrac{1}{3}, \dots\}\times[0,1]\big)$$

Note that $\mathcal{X}$ is locally compact and geodesic.

Its completion $\bar{\mathcal{X}}$ is isometric to the closure of $\mathcal{X}$ equipped with the induced length metric; $\bar{\mathcal{X}}$ is obtained from $\mathcal{X}$ by adding two points $p=(0,0)$ and $q=(0,1)$.

The point p admits no compact neighborhood in $\bar{\mathcal{X}}$ and there is no geodesic connecting p to q in $\bar{\mathcal{X}}$.

Comment. This example is taken from the book by Martin Bridson and André Haefliger [25, I.3.6(4)].

1.8. Let $\mathcal{W}=\mathcal{U}\times\mathcal{V}$. Choose two pairs of points $u^0,u^1\in\mathcal{U}$ and $v^0,v^1\in\mathcal{V}$. Set $a=|u^0-u^1|_{\mathcal{U}}$, $b=|v^0-v^1|_{\mathcal{V}}$ and $c=|(u^0,v^0)-(v^1,u^1)|_{\mathcal{W}}$.

Since $\mathcal{U}$ and $\mathcal{V}$ are length spaces, given $\varepsilon>0$, we can choose curves $\alpha\colon[0,1]\to\mathcal{U}$ from u^0 to u^1 and $\beta\colon[0,1]\to\mathcal{V}$ from v^0 to v^1 such that

$$\operatorname{length}\alpha<a+\varepsilon \quad\text{and}\quad \operatorname{length}\beta<b+\varepsilon.$$

Reparametrizing the paths proportional their lengths we can assume that α is $(a+\varepsilon)$-Lipschitz and β is $(b+\varepsilon)$-Lipschitz. Therefore the path $t\mapsto(\alpha(t),\beta(t))$ is $(c+2\cdot\varepsilon)$-Lipschitz. In particular, its length cannot exceed $c+2\cdot\varepsilon$ for any $\varepsilon>0$. Hence $\mathcal{W}$ meets the conditions in the definition of length space.

1.9. Let $\gamma\colon t\mapsto(\alpha(t),\beta(t))$ be a geodesic path in $\mathcal{W}=\mathcal{U}\times\mathcal{V}$. Show and use that

$$\begin{aligned}|\alpha(t_0)-\alpha(t_1)|_{\mathcal{U}}&=a\cdot|\gamma(t_0)-\gamma(t_1)|_{\mathcal{W}}\\|\beta(t_0)-\beta(t_1)|_{\mathcal{U}}&=b\cdot|\gamma(t_0)-\gamma(t_1)|_{\mathcal{W}}\end{aligned}$$

for any t_0 and t_1 and some fixed values $a\geqslant 0$ and $b\geqslant 0$ such that $a^2+b^2=1$.

1.10. Let γ be a unit-speed parametrization of $[pq]$. Show that after shifting the parametrization, we can assume that $|\gamma(t)|=\sqrt{a^2+t^2}$ for some constant a.

Let $\hat\gamma(t)$ be the projection of $\gamma(t)$ to $\mathcal{U}$. Show and use that $t\mapsto\hat\gamma(a\cdot\tan t)$ is a geodesic in $\mathcal{U}$.

1.11. A point in $\mathbb{R}\times\operatorname{Cone}\mathcal{U}$ can be described by a triple (x,r,p), where $x\in\mathbb{R}$, $r\in\mathbb{R}_{\geqslant}$ and $p\in\mathcal{U}$. Correspondingly, a point in $\operatorname{Cone}[\operatorname{Susp}\mathcal{U}]$ can be described by a triple (ρ,φ,p), where $\rho\in\mathbb{R}_{\geqslant}$, $\varphi\in[0,\pi]$ and $p\in\mathcal{U}$.

The map $\operatorname{Cone}[\operatorname{Susp}\mathcal{U}]\to\mathbb{R}\times\operatorname{Cone}\mathcal{U}$ defined as

$$(\rho,\varphi,p)\mapsto(\rho\cdot\cos\varphi,\rho\cdot\sin\varphi,p)$$

is the needed isometry.

1.15. Assume the contrary; that is

$$\measuredangle[p\,{}^x_z] + \measuredangle[p\,{}^y_z] < \pi.$$

By the triangle inequality for angles (1.14) we have

$$\measuredangle[p\,{}^x_y] < \pi.$$

The latter contradicts the triangle inequality for the triangle $[\bar{x}p\bar{y}]$, where the points $\bar{x} \in \,]px]$ and $\bar{y} \in \,]py]$ are sufficiently close to p.

1.18. By definition of Hausdorff convergence

$$p \in A_\infty \qquad \Longleftrightarrow \qquad \mathrm{dist}_{A_n}(p) \to 0 \quad \text{as} \quad n \to \infty.$$

The latter is equivalent to the existence of a sequence $p_n \in A_n$ such that $|p_n - p| \to 0$ as $n \to \infty$; or equivalently $p_n \to p$. Hence the first statement follows.

The converse is false. For example, consider the alternating sequence of two distinct closed sets $A, B, A, B, \dots$; note that it is not a converging sequence in the sense of Hausdorff. On the other hand, the set of all limit points is well defined — it is the intersection $A \cap B$.

Comment. The set $\underline{A}_\infty$ of all limits of sequences $p_n \in A_n$ is called the lower closed limit and the set $\bar{A}_\infty$ of all partial limits of such sequences is called the upper closed limit. Clearly $\underline{A}_\infty \subset \bar{A}_\infty$. If $\underline{A}_\infty = \bar{A}_\infty$, then it is called the closed limit of A_n. All these convergences were introduced by Felix Hausdorff in [54].

For the class of closed subsets of a proper metric spaces, closed limits coincide with limits in the sense of Hausdorff as we defined them.

1.23. Given any pair of points $x_0, y_0 \in \mathcal{K}$, consider two sequences (x_n) and (y_n) such that $x_{n+1} = f(x_n)$ and $y_{n+1} = f(y_n)$ for each n.

Since $\mathcal{K}$ is compact, we can choose an increasing sequence of integers n_i such that both sequences $(x_{n_i})_{i=1}^\infty$ and $(y_{n_i})_{i=1}^\infty$ converge. In particular, both of these sequences are Cauchy; that is,

$$|x_{n_i} - x_{n_j}|_{\mathcal{K}}, |y_{n_i} - y_{n_j}|_{\mathcal{K}} \to 0 \quad \text{as} \quad \min\{i, j\} \to \infty.$$

Since f is distance non-decreasing, we get

$$|x_0 - x_{|n_i - n_j|}| \leqslant |x_{n_i} - x_{n_j}|.$$

It follows that there is a sequence $m_i \to \infty$ such that

$$(*) \qquad x_{m_i} \to x_0 \quad \text{and} \quad y_{m_i} \to y_0 \quad \text{as} \quad i \to \infty.$$

Set

$$\ell_n = |x_n - y_n|_{\mathcal{K}}.$$

Since f is distance non-decreasing, (ℓ_n) is a non-decreasing sequence.

By $(*)$, $\ell_{m_i} \to \ell_0$ as $m_i \to \infty$. It follows that (ℓ_n) is a constant sequence.

In particular

$$|x_0 - y_0|_{\mathcal{K}} = \ell_0 = \ell_1 = |f(x_0) - f(y_0)|_{\mathcal{K}}$$

for any pair x_0 and y_0. That is, f is distance-preserving, in particular, injective.

From $(*)$, we also get that $f(\mathcal{K})$ is everywhere dense. Since $\mathcal{K}$ is compact, $f\colon \mathcal{K} \to \mathcal{K}$ is surjective. Hence the result.

Comment. This exercise is a basic introductory lemma on Gromov–Hausdorff distance (see for example [27, 7.3.30]). The presented proof is not quite standard, it was found by Travis Morrison.

1.24. To prove part *(a)*, fix a countable dense set of points $\mathfrak{S} \subset \mathcal{X}_\infty$. For each point $x \in \mathfrak{S}$, choose a sequence of points $x_n \in \mathcal{X}_n$ such that $x_n \xrightarrow{\rho} x$.

Applying the diagonal procedure, we can pass to a subsequence of $\mathcal{X}_n$ such that each of the constructed sequences ρ'-converge; that is, $x_n \xrightarrow{\rho'} x'$ for some $x' \in \mathcal{X}'_\infty$.

In this way we get a map $\mathfrak{S} \to \mathcal{X}'_\infty$ defined as $x \mapsto x'$. Note that this map preserves distances and therefore can be extended to a distance-preserving map $\mathcal{X}_\infty \to \mathcal{X}'_\infty$. Likewise we construct a distance-preserving map $\mathcal{X}'_\infty \to \mathcal{X}_\infty$.

It remains to apply Exercise 1.23.

The proof of part *(b)* is nearly identical, but one has to apply Exercise 1.23 to closed balls centered at the limits of x_n in $\mathcal{X}_\infty$ and $\mathcal{X}'_\infty$.

2.3. Note that it is sufficient to show that $\tilde\measuredangle(p\,{}^{\bar x}_{y}) \leqslant \tilde\measuredangle(p\,{}^{x}_{y})$ for any $\bar x \in \,]px[$. The latter follows from Alexandrov's lemma (1.12) and the CAT(0) comparison for the quadruple $p, x, \bar x, y$.

2.4. Observe that $|h_t(x) - h_t(y)| \leqslant t\cdot|x - y|$ and therefore

$$|h_{t_0}(x) - h_{t_1}(y)| \leqslant t_0\cdot|x - y| + |t_0 - t_1|\cdot|p - y|.$$

Make a conclusion.

2.5. Assume that a geodesic $[px]$ cannot be extended behind x. Apply the homotopy from 2.4 to prove that $\mathcal{U}$ has vanishing local homology (and/or homotopy) groups of at x. Use that manifolds have some nontrivial local homologies (and/or homotopy) groups.

2.6. Fix a sufficiently small $\varepsilon > 0$. Recall that by Proposition 2.9, any local geodesic in $\mathcal{U}$ is a geodesic.

Consider a sequence of directions ξ_n at p of geodesics $[pq_n]$. We can assume that the distances $|p - q_n|_{\mathcal{U}}$ are equal to ε for all n; here we use that the geodesics are extendable.

Since $\mathcal{U}$ is proper, we can pass to a converging subsequence of q_n; suppose q is its limit. Show that the direction ξ of $[pq]$ is the limit of directions ξ_n.

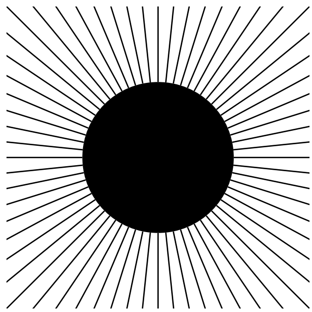

The unit disc in the plane with attached half-line to each point of its boundary is a complete CAT(0) length space with extendable geodesics. However, the space of geodesic directions on the boundary of the disc is not complete — there is no geodesic tangent to the boundary of the disc. This provides a counterexample to the statement of the exercise if $\mathcal{U}$ is not assumed to be proper.

2.10. It is sufficient to prove Jensen's inequality

$$h(\tfrac{t_0+t_1}{2}) \leqslant \tfrac{1}{2}(h(t_0) + h(t_1))$$

where $h(t) := f \circ \gamma(t) - \frac{1}{2}{\cdot}t^2$. Observe that the inequality holds in Euclidean plane, and apply that triangle $[p\,\gamma(t_0)\,\gamma(t_1)]_{\mathcal{U}}$ is thin.

2.11. Observe that it is sufficient to show that

$$|\gamma_1(t) - \gamma_2(t)|_{\mathcal{U}} \leqslant (1-t){\cdot}|\gamma_1(0) - \gamma_2(0)|_{\mathcal{U}} + t{\cdot}|\gamma_1(1) - \gamma_2(1)|_{\mathcal{U}}.$$

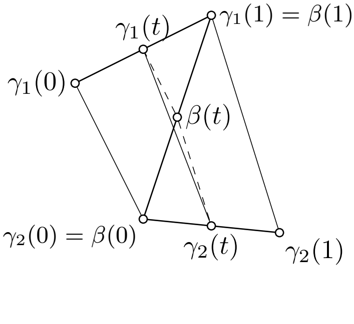

Let β be the geodesic path from $\gamma_1(0)$ to $\gamma_2(1)$. Observe that

$$\begin{aligned}|\gamma_1(t)-\beta(t)|_{\mathcal{U}} &\leqslant t\cdot|\gamma_1(1)-\beta(1)|_{\mathcal{U}},\\ |\beta(t)-\gamma_2(t)|_{\mathcal{U}} &\leqslant (1-t)\cdot|\beta(0)-\gamma_2(0)|_{\mathcal{U}},\end{aligned}$$

and apply the triangle inequality.

2.12. It is sufficient to show that

$$(*)\qquad \operatorname{dist}_A\circ\gamma(t)\leqslant(1-t)\cdot\operatorname{dist}_A\circ\gamma(0)+t\cdot\operatorname{dist}_A\circ\gamma(1).$$

for any geodesic path γ. Note that given $\varepsilon>0$, there are points $p,q\in A$ such that

$$|p-\gamma(0)|<\operatorname{dist}_A\circ\gamma(0)+\varepsilon\quad\text{and}\quad|q-\gamma(1)|<\operatorname{dist}_A\circ\gamma(1)+\varepsilon.$$

Let β be a geodesic path from p to q. By 2.11,

$$|\beta(t)-\gamma(t)|\leqslant(1-t)\cdot|p-\gamma(0)|+t\cdot|q-\gamma(1)|.$$

Since A is convex, $\beta(t)\in A$ for any t. Since $\varepsilon>0$ is arbitrary, we get $(*)$.

2.13. Since $\mathcal{U}$ is proper, the set $K\cap\overline{\mathrm{B}}[p,R]$ is compact for any $R<\infty$. The existence of at least one point p^* that minimizes the distance from p follows.

Assume p^* is not uniquely defined; that is, two distinct points in K, say x and y, minimize the distance from p. Since K is convex, the midpoint z of $[xy]$ lies in K.

Thinness of triangles implies that

$$|p-z|<|p-x|=|p-y|,$$

a contradiction.

It remains to show that the map $p\mapsto p^*$ is short, that is,

$$(*)\qquad |p-q|\geqslant|p^*-q^*|.$$

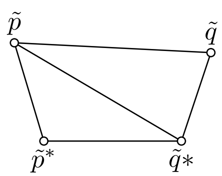

for any $p,q\in\mathcal{U}$.

Assume $p\neq p^*$, $q\neq q^*$, $p^*\neq q^*$. Construct the model triangles $[\tilde p\tilde p^*\tilde q^*]$ and $[\tilde p\tilde q\tilde q^*]$ of $[pp^*q^*]$ and $[pqq^*]$ so that the points $\tilde p^*$ and $\tilde q$ lie on the opposite sides from $[\tilde p\tilde q^*]$.

From thinness of triangles $[pp^*q^*]$ and $[pqq^*]$, we get that

$$\measuredangle[\tilde p^*\,{}^{\tilde p}_{\tilde q^*}],\measuredangle[\tilde q^*\,{}^{\tilde p^*}_{\tilde q}]\geqslant\tfrac{\pi}{2}.$$

Hence

$$|\tilde{p} - \tilde{q}| \geqslant |\tilde{p}^* - \tilde{q}^*|.$$

The latter is equivalent to $(*)$.

In the remaining cases: $(*)$ holds automatically if (1) $p^* = q^*$ or (2) $p = p^*$ and $q = q^*$. If $p = p^*$, $q \neq q^*$ and $p^* \neq q^*$, then thinness of $[pqq^*]$ implies that

$$\measuredangle[\tilde{q}^*\,{}^{\tilde{q}}_{\tilde{p}}] \geqslant \tfrac{\pi}{2},$$

and $(*)$ follows.

Comment. It is sufficient to assume that the space is complete length and CAT(0); see [9].

2.14. Fix a closed, connected, locally convex set K. Apply 2.12 to show that dist_K is convex in a neighborhood $\Omega \supset K$; that is, dist_K is convex along any geodesic completely contained in Ω.

Since K is locally convex, it is locally path connected. Since K is connected, it is also path connected.

Fix two points $x, y \in K$. Let us connect x to y by a path $\alpha\colon [0,1] \to K$. Use 2.8 to show that the geodesic $[x\,\alpha(s)]$ is uniquely defined and depends continuously on s.

If $[xy] = [x\,\alpha(1)]$ does not completely lie in K, then there is a value $s \in [0,1]$ such that $[x\,\alpha(s)]$ lies in Ω, but does not completely lie in K. Therefore $f = \operatorname{dist}_K$ is convex along $[x\alpha(s)]$.

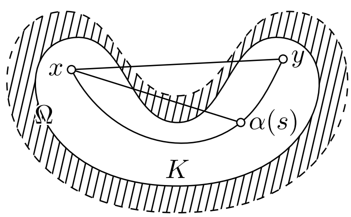

Note that $f(x) = f(\alpha(s)) = 0$ and $f \geqslant 0$, therefore $f(z) = 0$ for any $z \in [x\,\alpha(s)]$; that is, $[x\,\alpha(s)] \subset K$, a contradiction.

Comment. This statement generalizes the theorem of Heinrich Tietze, and the proof presented here is nearly identical to the original proof given in [88].

2.18. The "if" part follows from Reshetnyak gluing theorem (2.16).

Assume $\mathcal{W}$ is CAT(0). Note that one copy of $\mathcal{U}$ embeds isometrically in $\mathcal{W}$. Conclude that $\mathcal{U}$ is CAT(0).

Assume A is not convex; that is $[xy] \not\subset A$ for some $x, y \in A$. Observe that there are distinct geodesics from x to y in $\mathcal{W}$. Arrive at a contradiction with the uniqueness of geodesics (2.2).

3.7. By approximation, it is sufficient to consider the case when A and B have smooth boundary.

If $[xy] \cap A \cap B \neq \varnothing$, then $z_0 \in [xy]$ and $\dot{A}, \dot{B}$ can be chosen to be arbitrary half-spaces containing A and B respectively.

In the remaining case $[xy] \cap A \cap B = \varnothing$, we have $z_0 \in \partial(A \cap B)$. Consider the solid ellipsoid

$$C = \{\, z \in \mathbb{E}^m \;:\; f(z) \leqslant f(z_0) \,\}.$$

Note that C is compact, convex and has smooth boundary.

Suppose $z_0 \in \partial A \cap \operatorname{Int} B$. Then A and C touch at z_0 and we can set $\dot{A}$ to be the uniquely defined supporting half-space to A at z_0 and $\dot{B}$ to be any half-space containing B. The case $z_0 \in \partial B \cap \operatorname{Int} A$ is treated similarly.

Finally, suppose $z_0 \in \partial A \cap \partial B$. Then the set $\dot{A}$ (respectively, $\dot{B}$) is defined as the unique supporting half-space to A (respectively, B) at z_0 containing A (respectively, B).

Suppose $f(z) < f(z_0)$ for some $z \in \dot{A} \cap \dot{B}$. Since f is concave, $f(\bar{z}) < f(z_0)$ for any $\bar{z} \in [zz_0[$. Since $[zz_0[\cap A \cap B \neq \varnothing$, the latter contradicts the fact that z_0 is minimum point of f on $A \cap B$.

3.8. Fix two open balls $B_1 = \mathrm{B}(0, r_1)$ and $B_2 = \mathrm{B}(0, r_2)$ such that

$$B_1 \subset A^i \subset B_2$$

for each wall A^i.

Note that all the intersections of the walls have ε-wide corners for

$$\varepsilon = 2 \cdot \arcsin \tfrac{r_1}{r_2}.$$

The proof can be guessed from the picture.

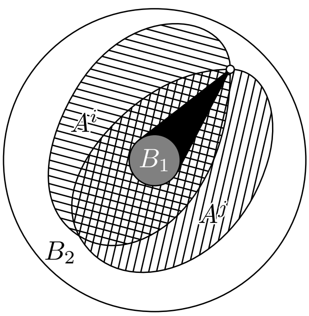

3.9. Note that any centrally symmetric convex closed set in Euclidean space is a product of a compact centrally symmetric convex set and a subspace.

It follows that there is $R < \infty$ such that if X is an intersection of an arbitrary number of walls, then for any point $p \in X$ there is an isometry of X that moves p to a point in the ball $\mathrm{B}(0, R)$.

It remains to repeat the proof of Exercise 3.8.

3.13. Imagine that each ball has zero radius, then may think that balls pass thru each other. That is, every ball moves with constant speed along the line. Let $x_i(t)$ be the coordinate of i^{th} ball at time t. Note that the graph x_i in the (t, x)-plane is a straight line. Every two lines have at most one intersection and each collision corresponds to one such intersection. We have n lines and therefore at most $\binom{n}{2}$ collisions.

The general case, with balls of radius $r > 0$ can be reduced to the case above. To do this, one have to exclude the space occupied by the balls. In other words, if the i^{th} ball centered at x, then the we assume that its coordinate is $x - 2{\cdot}r{\cdot}i$. With these new coordinates, the balls behave exactly as the balls with vanishing radii.

3.14. Assume that trajectory γ defined in an interval $[a, b)$ and collisions accumulate at b. Consider the infinite puff pastry $\mathcal{R}_\gamma$ for γ and let $\bar\gamma$ be its lift.

Pass to the completion $\bar{\mathcal{R}}_\gamma$ of $\mathcal{R}_\gamma$; observe that $\bar{\mathcal{R}}_\gamma$ is CAT(0). Define $\bar\gamma(b)$ to be the limit point of $\bar\gamma(t)$ as $t \to b$. Notice that $\bar\gamma\colon [a, b] \to \bar{\mathcal{R}}_\gamma$ is a minimizing geodesic.

Show that $\bar\gamma(b)$ lies in the lift of the intersection of all walls; in other words, $\bar\gamma(b)$ belongs to the intersection of all the levels of the puff pastry. Conclude that $\bar\gamma$ completely lies on the lowest level, and arrive at a contradiction.

Finally, by taking two tangent discs as the walls of a billiard table, we obtain the needed example. Indeed, a trajectory that starts near the common point of the discs in the direction perpendicular to their common tangent pane will have to bounce intensively for quite a while.

4.9. Apply 2.10.

4.11. Suppose D is a convex figure that majorizes $[p_1 \dots p_n]$. Show that D is an n-gon and, for each i, the majorization sends to p_i its vertex, say $\tilde p_i$. Show that the external angle of D at $\tilde p_i$ cannot exceed the exteral angle at p_i; make a conclusion.

4.12. It is a theorem of the first author and Richard Bishop [6]. The required polygon is shown on the diagram; it lies in the product space of the real line and a tripod; that is, three line segments glued together at one end. Note that in the original Fáry–Milnor theorem, the inequality is strict.

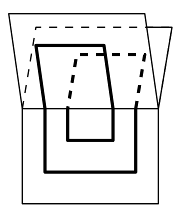

4.13. Apply the majorization theorem and the standard arm lemma.

5.5. Note that the existence of a null-homotopy is equivalent to the following. There are two one-parameter families of paths α_τ and β_τ, $\tau \in [0, 1]$ such that:

- ◇ $\operatorname{length}\alpha_\tau, \operatorname{length}\beta_\tau < \pi$ for any τ.
- ◇ $\alpha_\tau(0) = \beta_\tau(0)$ and $\alpha_\tau(1) = \beta_\tau(1)$ for any τ.
- ◇ $\alpha_0(t) = \beta_0(t)$ for any t.

◇ $\alpha_1(t) = \alpha(t)$ and $\beta_1(t) = \beta(t)$ for any t.

By Corollary 5.3, the construction in Corollary 5.4 produces the same result for α_τ and β_τ. Hence the result.

5.9. By the globalization theorem there is a nontrivial homotopy class of closed curves.

Consider a shortest noncontractible closed curve γ in $\mathcal{X}$; note that such a curve exists.

Indeed, let ℓ be the infimum of lengths of all noncontractible closed curves in $\mathcal{X}$. Geodesic homotopy construction implies that for two sufficiently close closed curves in $\mathcal{X}$ are homotopic. Then choosing a sequence of unit speed noncontractible curves whose lengths converge to ℓ, an Arzelá–Ascoli type of argument shows that these curves subconverge to a noncontractible curve γ of length ℓ.

Assume that γ is not a geodesic circle, that is, there are two points p and q on γ such that the distance $|p-q|$ is shorter then the lengths of the arcs, say α_1 and α_2, of γ from p to q. Consider the products, say γ_1 and γ_2, of $[qp]$ with α_1 and α_2. Then γ_1 or γ_2 is noncontractible,

$$\operatorname{length}\gamma_1 < \ell \quad \text{and} \quad \operatorname{length}\gamma_2 < \ell$$

— a contradiction.

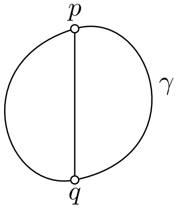

In the CAT(1) case we also have a geodesic circle. The proof is done nearly the same way, but we need to consider the homotopy classes of closed curves shorter than $2\cdot\pi$. One also need to apply 5.5, to show that curves γ_1 and γ_2 are not contractible in the class of curves shorter than $2\cdot\pi$.

Comment. Note that the surface of revolution of the graph of $y = e^x$ around the x-axis is locally CAT(0) but has no closed geodesics. Therefore, in this exercise, one cannot trade compactness to properness.

5.10. Consider a closed ε-neighborhood A of the geodesic. Note that A_ε is convex. By the Reshetnyak gluing theorem, the double $\mathcal{W}_\varepsilon$ of $\mathcal{U}$ along A_ε is CAT(0).

Consider the space $\mathcal{W}'_\varepsilon$ obtained by doubly covering $\mathcal{U} \setminus A_\varepsilon$ and gluing back A_ε.

Observe that $\mathcal{W}'_\varepsilon$ is locally isometric to $\mathcal{W}_\varepsilon$. That is, for any point $p' \in \mathcal{W}'_\varepsilon$ there is a point $p \in \mathcal{W}_\varepsilon$ such that the δ-neighborhood of p' is isometric to the δ-neighborhood of p for all small $\delta > 0$.

Further observe that $\mathcal{W}'_\varepsilon$ is simply connected since it admits a deformation retraction onto A_ε, which is contractible. By the globalization theorem, $\mathcal{W}'_\varepsilon$ is CAT(0).

It remains to note that $\tilde{\mathcal{U}}$ can be obtained as a limit of $\mathcal{W}'_\varepsilon$ as $\varepsilon \to 0$, and apply 2.1.

6.6. Assume $\mathcal{P}$ is not CAT(0). Then by 6.5, a link Σ of some simplex contains a closed geodesic α with length $4\cdot\ell < 2\cdot\pi$. We can assume that Σ has minimal possible dimension; so, by 6.5, Σ is locally CAT(1).

Divide α into two equal arcs α_1 and α_2.

Assume α_1 and α_2 are length minimizing; parameterize them by $[-\ell, \ell]$. Fix a small $\delta > 0$ and consider two curves in $\operatorname{Cone}\Sigma$ written in polar coordinates as

$$\gamma_i(t) = (\alpha_i(\arctan \tfrac{t}{\delta}), \sqrt{\delta^2 + t^2}).$$

Observe that both curves γ_1 and γ_2 are geodesics in $\operatorname{Cone}\Sigma$ and have common ends.

Note that a small neighborhood of the tip of $\operatorname{Cone}\Sigma$ admits an isometric embedding into $\mathcal{P}$. Hence we can construct two geodesics γ_1 and γ_2 in $\mathcal{P}$ with common endpoints.

It remains to consider the case when α_1 (and therefore α_2) is not length minimizing.

Pass to its maximal length minimizing arc $\bar{\alpha}_1$ of α_1. Since Σ is locally CAT(1), 5.3 implies that there is another geodesic $\bar{\alpha}_2$ in Σ_p that shares endpoints with $\bar{\alpha}_1$. It remains to repeat the above construction for the pair $\bar{\alpha}_1$, $\bar{\alpha}_2$.

Comments. By 2.2, the given condition is a necessary and sufficient.

In the proof, one can apply 5.9; in this case the last part of the arguing is not needed.

6.7. Note that it is sufficient to construct a polyhedral space $\mathcal{P}$ homeomorphic to the 3-disc such that (1) $\mathcal{P}$ is locally CAT(0) in its interior and (2) the boundary of $\mathcal{P}$ is locally concave; in particular, each edge on the boundary of $\mathcal{P}$ has angle at least π.

Indeed, once $\mathcal{P}$ is constructed, taking the double of $\mathcal{P}$ along its boundary produces the needed metric on $\mathbb{S}^3$.

The construction of $\mathcal{P}$ goes along the same lines as the construction of a Riemannian metric on the 3-disc with concave boundary and negative sectional curvature. This construction is given by Joel Hass in [52].

Comments. By the globalization theorem (5.6) the obtained metric on $\mathbb{S}^3$ is not locally CAT(0).

This problem originated from a discussion in Oberwolfach between Brian Bowditch, Tadeusz Januszkiewicz, Dmitri Panov and the third author. Another solution was given by Karim Adiprasito [2]; he proved that an example can be found among spaces that admit a cubulation into unit cubes.

6.9. Checking the flag condition is straightforward once we know the following description of the barycentric subdivision.

Each vertex v of the barycentric subdivision corresponds to a simplex $\triangle_v$ of the original triangulation. A set of vertices forms a simplex in the subdivision if it can be ordered, say as $v_1, \dots, v_k$, so that the corresponding simplices form a nested sequence

$$\triangle_{v_1} \subset \dots \subset \triangle_{v_k}.$$

Comment. There are a compact metrizable contractible and locally contractible spaces that do not admit a CAT(1) length metrics [1].

6.13. Use induction on the dimension to prove that if in a spherical simplex $\triangle$ every edge is at least $\frac{\pi}{2}$, then all dihedral angles of $\triangle$ are at least $\frac{\pi}{2}$.

The rest of the proof goes along the same lines as the proof of the flag condition (6.12). The only difference is geodesic spends *at least* π on each visit to Star_v.

Comment. Note that it is not sufficient to assume only that the all dihedral angles of the simplices are at least $\frac{\pi}{2}$. Indeed, the two-dimensional sphere with removed interior of a small rhombus is a spherical polyhedral space glued from four triangles with all the angles at least $\frac{\pi}{2}$. On the other hand the boundary of the rhombus is closed local geodesic in this space. Therefore the space cannot be CAT(1).

6.14. Apply the globalization theorem (5.6) with 6.2 and 6.13.

6.15. The space $\mathcal{T}_n$ has a natural cone structure with the vertex formed by the completely degenerate tree — all its edges have zero length. Note that the space Σ over which the cone is taken comes naturally with a triangulation with all-right spherical simplicies.

The link of any simplex of this triangulation satisfies the no-triangle condition (6.8). Indeed, fix a simplex $\triangle$ of the complex; it can be described by combinatorics of a possibly degenerate tree. A triangle in the link of $\triangle$ can be described by three ways to resolve a degeneracy by adding one edge of positive length, such that (1) any pair of these resolutions can be done simultaneously, but (2) all three cannot be done simultaneously. Direct inspection shows that this is impossible.

Therefore, by Proposition 6.10 our complex is flag. It remains to apply the flag condition (6.12) and then Exercise 6.2.

6.17. If the complex $\mathcal{S}$ is flag, then its cubical analog $\square_{\mathcal{S}}$ is locally CAT(0) and therefore aspherical.

Assume now that the complex $\mathcal{S}$ is not flag. Extend it to a flag complex $\mathcal{T}$ by gluing a simplex in every clique (that is, a complete subgraph) of its one-skeleton.

Note that the cubical analog $\square_{\mathcal{S}}$ is a proper subcomplex in $\square_{\mathcal{T}}$. Since $\mathcal{T}$ is flag, $\tilde{\square}_{\mathcal{T}}$, the universal cover of $\square_{\mathcal{T}}$, is CAT(0). Let $\tilde{\square}_{\mathcal{S}}$ be the inverse image of $\square_{\mathcal{S}}$ in $\tilde{\square}_{\mathcal{T}}$.

Choose a cube Q with minimal dimension in $\tilde{\square}_{\mathcal{T}}$ which is not present in $\tilde{\square}_{\mathcal{S}}$. By Exercise 2.14, Q is a convex set in $\tilde{\square}_{\mathcal{T}}$. The closest point projection $\tilde{\square}_{\mathcal{T}} \to Q$ is a retraction. It follows that the boundary ∂Q is not contractible in $\tilde{\square}_{\mathcal{T}} \setminus \operatorname{Int} Q$. Therefore the spheroid ∂Q is not contractible in $\tilde{\square}_{\mathcal{S}}$. That is, a covering of $\square_{\mathcal{S}}$ is not aspherical and therefore $\square_{\mathcal{S}}$ is not as well.

6.20. The solution goes along the same lines as the proof of Lemma 6.19, but few changes are needed.

The cycle γ is taken in the complement $\mathcal{S} \setminus \{v\}$ (or, alternatively, in the link of v in $\mathcal{S}$). Instead of a vertex, one has to take edge e in $\tilde{Q}$ that corresponds to v; so we have to show existence of large cycle in $\tilde{Q}$ that is not contractible in $\tilde{Q} \setminus e$. Think that G is made from the squares parallel to the squares of the cubical complex which meet the edges of the complex orthogonally at their midpoints. Note that formally speaking, G is not a subcomplex of the cubical analog.

6.22. In the proof we apply the following lemma. It follows from the disjoint discs property; see [39, 44].

Lemma. *Let $\mathcal{S}$ be a finite simplicial complex that is homeomorphic to an m-dimensional homology manifold for some $m \geqslant 5$. Assume that all vertices of $\mathcal{S}$ have simply connected links. Then $\mathcal{S}$ is a topological manifold.*

Note that it is sufficient to construct a simplicial complex $\mathcal{S}$ such that

- ◇ $\mathcal{S}$ is a closed $(m-1)$-dimensional homology manifold;
- ◇ $\pi_1(\mathcal{S} \setminus \{v\}) \neq 0$ for some vertex v in $\mathcal{S}$;
- ◇ $\mathcal{S} \sim \mathbb{S}^{m-1}$; that is, $\mathcal{S}$ is homotopy equivalent to $\mathbb{S}^{m-1}$.

Indeed, assume such $\mathcal{S}$ is constructed. Then the suspension $\mathcal{R} = \operatorname{Susp} \mathcal{S}$ is an m-dimensional homology manifold with a natural triangulation coming from $\mathcal{S}$. According to the lemma, $\mathcal{R}$ is a topological manifold. According to the generalized Poincaré conjecture, $\mathcal{R} \simeq \mathbb{S}^m$; that is $\mathcal{R}$ is homeomorphic to $\mathbb{S}^m$. Since $\operatorname{Cone} \mathcal{S} \simeq \mathcal{R} \setminus \{s\}$ where s denotes a south pole of the suspension and $\mathbb{E}^m \simeq \mathbb{S}^m \setminus \{p\}$ for any point $p \in \mathbb{S}^m$, we get

$$\operatorname{Cone} \mathcal{S} \simeq \mathbb{E}^m.$$

It remains to construct $\mathcal{S}$. Fix an $(m-2)$-dimensional homology sphere Σ with a triangulation such that $\pi_1\Sigma \neq 0$. An example of that type exists for any $m \geqslant 5$; a proof is given in [58].

Remove from Σ the interior of one $(m-2)$-simplex. Denote the resulting complex by Σ'. Since $m \geqslant 5$, we have $\pi_1\Sigma = \pi_1\Sigma'$.

Consider the product $\Sigma' \times [0,1]$. Attach to it the cone over its boundary $\partial(\Sigma' \times [0,1])$. Denote by $\mathcal{S}$ the resulting simplicial complex and by v the tip of the attached cone.

Note that $\mathcal{S}$ is homotopy equivalent to the spherical suspension over Σ, which is a simply connected homology sphere and hence is homotopy equivalent to $\mathbb{S}^{m-1}$. Hence $\mathcal{S} \sim \mathbb{S}^{m-1}$.

The complement $\mathcal{S} \setminus \{v\}$ is homotopy equivalent to Σ'. Therefore

$$\pi_1(\mathcal{S} \setminus \{v\}) = \pi_1\Sigma' = \pi_1\Sigma \neq 0.$$

That is, $\mathcal{S}$ satisfies the conditions above.

6.24; *(b)* $\Rightarrow$ *(a).* Since any closed curve can be considered as a short map from a boundary of a disc with some metric, it can be extended to a short map from a disc. Therefore any injective space is simply connected.

Therefore the globalization theorem and flag condition (5.6 and 6.12) imply that it is sufficient to show that each link in Q is flag. Further, by 6.10 it is sufficient to show that link of each cube in Q satisfies no-triangle condition.

Arguing by contradiction, we can assume that no-triangle condition does not hold at a vertex v; that is, a zero-dimensional cube. In this case v is a vertex of there edges e_x, e_y, and e_z; each pair of edges belong to one of the squares s_x, s_y, and s_z with complementary index, but the squares s_x, s_y, s_z do not belong to one cube. For higher dimensional cubes we have a product of this configuration with a cube.

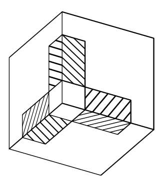

Let m_x, m_y and m_z be the midpoints of e_x, e_y, and e_z respectively. Consider 3 balls with centers m_x, m_y, and m_z and radius $\frac{1}{4}$. Observe that each pair of balls have a common point; but all three together have no points of intersection. It follows that (Q, ℓ^∞) is not an injective space — it does not contain a point on distance $\frac{1}{4}$ from m_x, m_y, and m_z — a contradiction.

(c) $\Rightarrow$ *(a).* Observe that median point $m(x,y,z)$ depends continuously on triple of points (x,y,z) and $m(x,x,y) = x$.

Given a loop $\gamma\colon [0,1] \to Q$ with base at $p = \gamma(0) = \gamma(1)$, consider the map $(a,b) \mapsto m(p, \gamma(a), \gamma(b))$ of the triangle $\triangle$ defined by $0 \leqslant a \leqslant$ $\leqslant b \leqslant 1$. Note that boundary of triangle runs along γ. It follows that γ is null homotopic and therefore Q is simply connected.

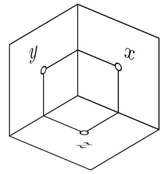

It remains to check that all links of Q satisfy no-triangle condition.

Assume that a link of Q does not satisfy the no-triangle condition. The same way as in the previous problem, we can assume that it is a link of a vertex; so we have a configuration of three squares s_x, s_y, and s_z, three edges e_x, e_y, and e_z, and one common vertex v as above. Observe that the centers x, y, and z of the squares s_x, s_y, and s_z. Observe that that the geodesics $[xy]_{\ell^1}$, $[xz]_{\ell^1}$, and $[yz]_{\ell^1}$ are uniquely defined and they have no common point. It follows that the triple (x,y,z) does not have a median; that is, (Q,ℓ^1) is not a median space — a contradiction.

7.12. Observe that the triangle $[pqx]$ is degenerate, in particular it is thin. It remains to apply the inheritance lemma (2.15).

7.13. By approximation, it is sufficient to consider the case when S has polygonal sides.

The latter case can be done by induction on the number of sides. The base case of triangle is evident.

To prove the induction step, apply Alexandrov's lemma (1.12) together with the construction in the inheritance lemma (2.15).

7.15. From Exercise 7.22, it follows that if a lune in $\mathbb{S}^2$ has perimeter smaller then $2\cdot\pi$, then it contains a closed hemisphere in its interior or lies in an open hemisphere. The same holds for a triangular region with concave sides.

By the assumption, Θ does not contain a closed hemisphere. That is, the first case cannot happen. It remains to apply the argument in the proof of Theorem 7.11.

7.18. The space $\tilde K$ is a cone over the branched covering Σ of $\mathbb{S}^3$ infinitely branching along two great circles.

If the planes are not orthogonal, then the minimal distance between the circles is less than $\frac{\pi}{2}$. Assume that this distance is realized by a geodesic $[\xi\zeta]$. The broken line made by four liftings of $[\xi\zeta]$ forms a closed local geodesic in Σ. By Proposition 2.9 (or Corollary 5.8), Σ is not CAT(1). Therefore by Exercise 6.2, K is not CAT(0).

If the planes are orthogonal, then the corresponding great circles in $\mathbb{S}^3$ are subcomplexes of a flag triangulation of $\mathbb{S}^3$ with all-right simplicies. The branching cover is also flag. It remains to apply the flag condition 6.12.

Comments. In [37], Ruth Charney and Michael Davis gave a complete answer to the analogous question for three planes. In particular they show that if a covering space of $\mathbb{E}^4$ branching at three planes intersecting at the origin is CAT(0), then these all are complex planes for some complex structure on $\mathbb{E}^4$.

7.22. Let α be a closed curve in $\mathbb{S}^2$ of length $2{\cdot}\ell$.

Assume $\ell < \pi$. Let $\check{\alpha}$ be a subarc of α of length ℓ, with endpoints p and q. Since $|p-q| \leqslant \ell < \pi$, there is a unique geodesic $[pq]$ in $\mathbb{S}^2$. Let z be the midpoint of $[pq]$.

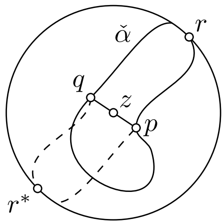

We claim that α lies in the open hemisphere H centered at z.

Assume the contrary; that is, α meets the equator ∂H at a point r. Without loss of generality we may assume that $r \in \check{\alpha}$.

The arc $\check{\alpha}$ together with its reflection in z form a closed curve of length $2{\cdot}\ell$ which meets r and its antipodal point r'. Thus $\ell = \operatorname{length}\check{\alpha} \geqslant |r - r'| = \pi$, a contradiction.

Solution via the Crofton formula. Let α be a closed curve in $\mathbb{S}^2$ of length $\leqslant 2{\cdot}\pi$. We wish to prove that α is contained in a hemisphere in $\mathbb{S}^2$. By approximation it suffices to prove this for smooth curves α of length $< 2{\cdot}\pi$ with transverse self-intersections.

Given $v \in \mathbb{S}^2$, denote by $v^\perp$ the equator in $\mathbb{S}^2$ with the pole at v. Further, $\#X$ will denote the number of points in the set X.

Obviously, if $\#(\alpha \cap v^\perp) = 0$, then α is contained in one of the hemispheres determined by $v^\perp$. Note that $\#(\alpha \cap v^\perp)$ is even for almost all v.

Therefore, if α does not lie in a hemisphere, then $\#(\alpha \cap v^\perp) \geqslant 2$ for almost all $v \in \mathbb{S}^2$.

By the Crofton formula we have that

$$\begin{aligned}\operatorname{length}(\alpha) &= \tfrac{1}{4}\cdot \int\limits_{v\in\mathbb{S}^2} \#(\alpha \cap v^\perp) \geqslant \\ &\geqslant 2{\cdot}\pi.\end{aligned}$$

7.23. Since Ω is not two-convex, we can fix a simple closed curve γ that lies in the intersection of a plane W_0 and Ω, and is contractible in Ω but not contractible in $\Omega \cap W_0$.

Let $\varphi\colon \mathbb{D} \to \Omega$ be a disc that shrinks γ. Applying the loop theorem (arguing as in the proof of Proposition 7.8), we can assume that φ is an embedding and $\varphi(\mathbb{D})$ lies on one side of W_0.

Let Q be the bounded closed domain cut from $\mathbb{E}^3$ by $\varphi(\mathbb{D})$ and W_0. By assumption it contains a point that is not in Ω. Changing W_0, γ

and φ slightly, we can assume that such a point lies in the interior of Q.

Fix a circle Γ in W_0 that surrounds $Q \cap W_0$. Since Q lies in a half-space with boundary W_0, there is a smallest spherical dome with boundary Γ that includes the set $R = Q \setminus \Omega$.

The dome has to touch R at some point p. The plane W tangent to the dome at p has the required property — the point p is an isolated point of the complement $W \setminus \Omega$. Further, by construction a small circle around p in W is contractible in Ω.

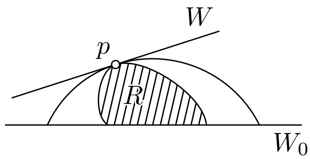

7.26. The proof is simple and visual, but it is hard to write it formally in a non-tedious way; for that reason we give only a sketch.

Consider the surface $\bar{S}$ formed by the closure of the remaining part S of the boundary. Note that the boundary ∂S of $\bar{S}$ is a collection of closed polygonal lines.

Assume $\bar{S}$ is not piecewise linear. Show that there is a line segment $[pq]$ in $\mathbb{E}^3$ that is tangent to $\bar{S}$ at some point p and has no common points with $\bar{S}$ except p.

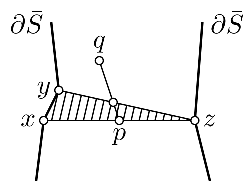

Since $\bar{S}$ is locally concave, there is a local inner supporting plane Π at p that contains the segment $[pq]$.

Note that $\Pi \cap \bar{S}$ contains a segment $[xy] \ni p$ with the ends in $\partial\bar{S}$. Denote by Π^+ the half-plane in Π that contains $[pq]$ and has $[xy]$ in its boundary.

Use the fact that $[pq]$ is tangent to S to show that there is a point $z \in \partial\bar{S}$ such that the line segment $[xz]$ or $[yz]$ lies in $\partial\bar{S} \cap \Pi^+$.

From the latter statement and local convexity of $\bar{S}$, it follows that the solid triangle $[xyz]$ lies in $\bar{S}$. In particular, all points on $[pq]$ sufficiently close to p lie in $\bar{S}$, a contradiction.

7.27. Suppose K is not two-convex; let γ be a closed simple curve in a plane W that do not meet Definition 7.2. Note that W is not a vertical plane; denote by V the three-dimensional subspace that spanned by W and vertical direction.

Note that γ is contractible in $V \cap K$. Act as in 7.23 to show that there exists a plane triangle $\triangle \subset V$ whose sides lie completely in K, but whose interior contains points from the complement $\mathbb{E}^m \setminus K$.

7.28. Clearly, the set W is two-convex, and so is K as the intersection of two-convex sets.

Consider two 2-dimensional hemispheres H_1 and H_2 in $\mathbb{S}^3$ such that the intersection $H_1 \cap H_2$ is a geodesic $[\xi\zeta]$ orthogonal to the boundary equators of H_1 and H_2 and

$$|\xi - \zeta|_{\mathbb{S}^3} < \tfrac{\pi}{2}.$$

Equip the complement $\mathbb{S}^3 \setminus (H_1 \cup H_2)$ with induced length metric and denote by Σ its completion.

Note that there is a closed geodesic in Σ whose projection to $\mathbb{S}^3$ is formed by a product of four copies of $[\xi\zeta]$. In particular there is a closed geodesic in Σ shorter than $2{\cdot}\pi$.

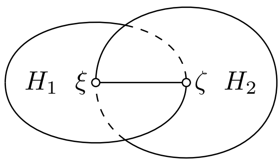

Hence Σ is not CAT(1) and therefore $K' = \operatorname{Cone}\Sigma$ is not CAT(0).

For a suitable choice of the motion ι, we have that $\tfrac{1}{n}{\cdot}K \to K'$ as $n \to \infty$ in the sense of Gromov–Hausdorff. By 2.1, K is not CAT(0).

8.3; *(a)* Suppose G is finite. Choose its orbit $\{p_1, \dots, p_n\}$. Consider the barycenter z of the array $\boldsymbol{p} = (p_1, \dots, p_n)$ with equal masses; in other words, $z := \mathfrak{S}_{\boldsymbol{p}}(\tfrac{1}{n}, \dots, \tfrac{1}{n})$. Observe that z is a fixed point of the acton.

(b). Let μ be the probability Haar measure on G. Choose a point $p \in \mathcal{U}$ and consider the function

$$f = \tfrac{1}{2}{\cdot} \int\limits_{g \in G} \operatorname{dist}^2_{g\cdot p}{\cdot}\mu.$$

Show that f is 1-convex. By 8.2, f has unique minimum point, say z. Observe that f if G-invariant; therefore, z is a fixed point.

8.6. The first part follows directly from the definitions. For the second part check a cone over circle with length bigger than $2{\cdot}\pi$, or a product of a tripod with the real line.

8.7; *(a).* Choose $q \in \bigcap_i B_i$. Assume $s \notin \bigcup_i B_i$. Observe that $|p_i - s| > |p_i - q|$ for any i and apply 8.5. Conclude that $s \notin \mathfrak{S}_{\boldsymbol{p}}(x)$.

(b). Show that there is a point, say z, that minimize

$$s = \max_i \{|p_i - z| - r_i\}.$$

Note that $s > 0$.

Let us show that $s = |p_i - z| - r_i$ for any i. Assume the contrary; that is, that $s > |p_j - z| - r_j$ for some j. Choose a point $q_j \in \bigcap_{i \neq j} B_i$.

Note that $s > |p_j - x| - r_j$ for any point $x \in]zq_j[$. Therefore s is not minimal — a contradiction.

Observe that $z \in \mathfrak{S}_{\boldsymbol{p}}(\triangle^k)$. By *(a)*, $\mathfrak{S}_{\boldsymbol{p}}(\partial\triangle^k) \subset \bigcup_i B_i$. Hence $\mathfrak{S}_{\boldsymbol{p}}(\triangle^k) \setminus \mathfrak{S}_{\boldsymbol{p}}(\partial\triangle^k) \neq \varnothing$.

(c). Choose $z \in \mathfrak{S}_{\boldsymbol{p}}(\triangle^k) \setminus \mathfrak{S}_{\boldsymbol{p}}(\partial\triangle^k)$. Show that for each j, there is a point z_j such that

$$|p_i - z_j| < |p_i - z|$$

for any $i \neq j$. Choose $r_i = \min_{j\neq i}\{|p_i - z_j|\}$.

8.16. Let k be the maximal integer such that any k subsets have a common points. By assumption $k \geqslant m + 1$.

Suppose $k < n$. We can assume that $n = k+1$ and $K_1 \cap \dots \cap K_n = \varnothing$. Choose a point array $\boldsymbol{p}$ such that

$$p_i \in \bigcap_{j\neq i} K_j$$

for each i. Observe that $\mathfrak{S}_{\boldsymbol{p}}(\partial\triangle^k) \subset \bigcup_i K_i$. Since $\mathfrak{S}_{\boldsymbol{p}}$ is degenerate, $\mathfrak{S}_{\boldsymbol{p}}(\triangle^k) \subset \bigcup_i K_i$. Apply Sperner's lemma to show that $\mathfrak{S}_{\boldsymbol{p}}(\triangle^k)(x) \in \bigcap_i K_i$ for some $x \in \triangle^k$, and arrive at a contradiction.

8.19. Let T be a binary rooted tree. Choose a metric on T such that each edge from level n to $n+1$ has length λ^n for for some $0 < \lambda < 1$. Passing to completion of the obtained space, we add to it a *crown* which is homeomorphic to the Cantor set. Observe that the completion has topological dimension 1. It remains to calculate the Hausdorff dimension of the crown for given λ.

Index

Bibliography

[1] A. Petrunin. *Topological spaces admitting CAT(1) metrics*. MathOverflow. eprint: `https://mathoverflow.net/q/460763`.

[2] K. Adiprasito. "A note on the simplex-cosimplex problem". *European J. Combin.* 66 (2017), 5–12.

[3] S. Alexander, D. Berg, and R. Bishop. "Geometric curvature bounds in Riemannian manifolds with boundary". *Trans. Amer. Math. Soc.* 339.2 (1993), 703–716.

[4] S. Alexander and R. Bishop. "Warped products of Hadamard spaces". *Manuscripta Math.* 96.4 (1998), 487–505.

[5] S. Alexander and R. Bishop. "The Hadamard–Cartan theorem in locally convex metric spaces". *Enseign. Math. (2)* 36.3-4 (1990), 309–320.

[6] S. Alexander and R. Bishop. "The Fary–Milnor theorem in Hadamard manifolds". *Proc. Amer. Math. Soc.* 126.11 (1998), 3427–3436.

[7] S. Alexander, V. Kapovitch, and A. Petrunin. "Alexandrov meets Kirszbraun". *Proceedings of the Gökova Geometry-Topology Conference 2010*. Int. Press, Somerville, MA, 2011, 88–109.

[8] S. Alexander, V. Kapovitch, and A. Petrunin. *An invitation to Alexandrov geometry: CAT(0) spaces*. 2019.

[9] S. Alexander, V. Kapovitch, and A. Petrunin. *Alexandrov geometry: foundations*. Vol. 236. Graduate Studies in Mathematics. 2024.

[10] A. D. Alexandrow. "Über eine Verallgemeinerung der Riemannschen Geometrie". *Schr. Forschungsinst. Math.* 1 (1957), 33–84.

[11] А. Д. Александров. «Линейчатые поверхности в метрических пространствах». *Вестник ЛГУ* 2 (1957), 15—44.

[12] A. Alexandroff. "The inner geometry of an arbitrary convex surface". *C. R. (Doklady) Acad. Sci. URSS (N.S.)* 32 (1941), 467–470.

[13] A. D. Alexandrow. "Über eine Verallgemeinerung der Riemannschen Geometrie". *Schr. Forschungsinst. Math.* 1 (1957), 33–84.

[14] Александр Данилович Александров. "Одна теорема о треугольниках в метрическом пространстве и некоторые ее приложения". *Труды МИАН СССР* 38.0 (1951), 5–23.

[15] F. Almgren. "Optimal isoperimetric inequalities". *Indiana Univ. Math. J.* 35.3 (1986), 451–547.

[16] F. D. Ancel, M. W. Davis, and C. R. Guilbault. “CAT(0) reflection manifolds”. *Geometric topology (Athens, GA, 1993).* Vol. 2. AMS/IP Stud. Adv. Math. Amer. Math. Soc., Providence, RI, 1997, 441–445.

[17] L. Andersson and R. Howard. “Comparison and rigidity theorems in semi-Riemannian geometry”. *Comm. Anal. Geom.* 6.4 (1998), 819–877.

[18] W. Ballmann. *Lectures on spaces of nonpositive curvature.* Vol. 25. DMV Seminar. With an appendix by Misha Brin. 1995.

[19] W. Ballmann. *Lectures on spaces of nonpositive curvature.* Vol. 25. DMV Seminar. With an appendix by Misha Brin. 1995.

[20] L. Billera, S. Holmes, and K. Vogtmann. “Geometry of the space of phylogenetic trees”. *Adv. in Appl. Math.* 27.4 (2001), 733–767.

[21] R. Bishop. “The intrinsic geometry of a Jordan domain”. *Int. Electron. J. Geom.* 1.2 (2008), 33–39.

[22] W. Blaschke. *Kreis und kugel.* Veit Leipzig, 1916.

[23] B. H. Bowditch. “Notes on locally CAT(1) spaces”. *Geometric group theory (Columbus, OH, 1992).* Vol. 3. Ohio State Univ. Math. Res. Inst. Publ. de Gruyter, Berlin, 1995, 1–48.

[24] B. Bowditch. “Median and injective metric spaces”. *Math. Proc. Cambridge Philos. Soc.* 168.1 (2020), 43–55.

[25] M. Bridson and A. Haefliger. *Metric spaces of non-positive curvature.* Vol. 319. Grundlehren der Mathematischen Wissenschaften. 1999.

[26] D. Burago. “Hard balls gas and Alexandrov spaces of curvature bounded above”. *Proceedings of the International Congress of Mathematicians, Vol. II (Berlin, 1998).* Extra Vol. II. 1998, 289–298.

[27] D. Burago, Y. Burago, and S. Ivanov. *A course in metric geometry.* Vol. 33. Graduate Studies in Mathematics. 2001. [Русский перевод: Бураго Д. Ю., Бураго Ю. Д., Иванов С. В. «Курс метрической геометрии», 2004.]

[28] D. Burago, S. Ferleger, and A. Kononenko. “Uniform estimates on the number of collisions in semi-dispersing billiards”. *Ann. of Math. (2)* 147.3 (1998), 695–708.

[29] D. Burago, S. Ferleger, and A. Kononenko. “Topological entropy of semi-dispersing billiards”. *Ergodic Theory Dynam. Systems* 18.4 (1998), 791–805.

[30] D. Burago, D. Grigoriev, and A. Slissenko. “Approximating shortest path for the skew lines problem in time doubly logarithmic in 1/epsilon”. *Theoret. Comput. Sci.* 315.2-3 (2004), 371–404.

[31] D. Burago and S. Ivanov. “Examples of exponentially many collisions in a hard ball system”. *Ergodic Theory Dynam. Systems* 41.9 (2021), 2754–2769.

[32] H. Busemann. “Spaces with non-positive curvature”. *Acta Math.* 80 (1948), 259–310.

[33] S. V. Buyalo. “Volume and fundamental group of a manifold of nonpositive curvature”. *Mat. Sb. (N.S.)* 122(164).2 (1983), 142–156.

[34] É. Cartan. *Leçons sur la Géométrie des Espaces de Riemann.* 1928.

[35] R. Charney and M. Davis. “Singular metrics of nonpositive curvature on branched covers of Riemannian manifolds”. *Amer. J. Math.* 115.5 (1993), 929–1009.

[36] R. Charney and M. Davis. “Strict hyperbolization”. *Topology* 34.2 (1995), 329–350.

[37] R. Charney and M. Davis. "Singular metrics of nonpositive curvature on branched covers of Riemannian manifolds". *Amer. J. Math.* 115.5 (1993), 929–1009.

[38] C. B. Croke. "A sharp four-dimensional isoperimetric inequality". *Comment. Math. Helv.* 59.2 (1984), 187–192.

[39] Robert J Daverman. *Decompositions of manifolds.* Academic Press, 1986.

[40] M. Davis. "Groups generated by reflections and aspherical manifolds not covered by Euclidean space". *Ann. of Math. (2)* 117.2 (1983), 293–324.

[41] M. Davis. "Exotic aspherical manifolds". *Topology of high-dimensional manifolds, No. 1, 2 (Trieste, 2001).* Vol. 9. ICTP Lect. Notes. Abdus Salam Int. Cent. Theoret. Phys., Trieste, 2002, 371–404.

[42] M. Davis and T. Januszkiewicz. "Hyperbolization of polyhedra". *J. Differential Geom.* 34.2 (1991), 347–388.

[43] M. Davis, T. Januszkiewicz, and J.-F. Lafont. "4-dimensional locally CAT(0)-manifolds with no Riemannian smoothings". *Duke Math. J.* 161.1 (2012), 1–28.

[44] R. Edwards. "The topology of manifolds and cell-like maps". *Proceedings of the International Congress of Mathematicians (Helsinki, 1978).* Citeseer. 1980, 111–127.

[45] R. Engelking. *Dimension theory.* 1978.

[46] Z. Frolík. "Concerning topological convergence of sets". *Czechoslovak Mathematical Journal* 10.2 (1960), 168–180.

[47] Г. А. Гальперин. «О системах локально взаимодействующих и отталкивающихся частиц, движущихся в пространстве». *Труды Московского математического общества* 43.0 (1981), 142—196.

[48] M. Gromov. "Hyperbolic groups". *Essays in group theory.* Vol. 8. Math. Sci. Res. Inst. Publ. Springer, New York, 1987, 75–263.

[49] M. Gromov. "Sign and geometric meaning of curvature". *Rend. Sem. Mat. Fis. Milano* 61 (1991), 9–123 (1994).

[50] M. Gromov. "Asymptotic invariants of infinite groups". *Geometric group theory, Vol. 2 (Sussex, 1991).* Vol. 182. London Math. Soc. Lecture Note Ser. 1993, 1–295.

[51] J. Hadamard. "Sur la forme des lignes géodésiques à l'infini et sur les géodésiques des surfaces réglées du second ordre". *Bull. Soc. Math. France* 26 (1898), 195–216.

[52] J. Hass. "Bounded 3-manifolds admit negatively curved metrics with concave boundary". *J. Differential Geom.* 40.3 (1994), 449–459.

[53] A. Hatcher. *Algebraic topology.* 2002. [Русский перевод: Хатчер А. «Алгебраическая топология», 2011.]

[54] F. Hausdorff. *Grundzüge der mengenlehre.* Vol. 7. von Veit, 1914.

[55] W. Hurewicz and H. Wallman. *Dimension Theory.* Princeton Mathematical Series, v. 4. 1941.

[56] S. Ivanov. "On Helly's theorem in geodesic spaces". *Electron. Res. Announc. Math. Sci.* 21 (2014), 109–112.

[57] С. В. Иванов. *Об ограничении числа отражений в бильярдах в нормированных пространствах.* Геометрический семинар им. А. Д. Александрова, ПОМИ. ауд. 203, 17:00, 22 апреля, 2024.

[58] M. Kervaire. “Smooth homology spheres and their fundamental groups”. *Trans. Amer. Math. Soc.* 144 (1969), 67–72.

[59] B. Kleiner. “The local structure of length spaces with curvature bounded above”. *Math. Z.* 231.3 (1999), 409–456.

[60] B. Kleiner. “An isoperimetric comparison theorem”. *Invent. Math.* 108.1 (1992), 37–47.

[61] U. Lang and V. Schroeder. “Kirszbraun’s theorem and metric spaces of bounded curvature”. *Geom. Funct. Anal.* 7.3 (1997), 535–560.

[62] Н. Лебедева и А. Петрунин. *Теорема Александрова о вложении многогранников.* 2022. [English translation: N. Lebedeva, A. Petrunin, *Alexandrov’s embedding theorem.*]

[63] И. М. Либерман. «Геодезические линии на выпуклых поверхностях». *Докл. АН СССР* 32.2 (1941), 310—312.

[64] A. Lytchak and S. Wagner. “Curvature bounds on length-minimizing discs”. *Geom. Dedicata* 218.2 (2024), Paper No. 49, 21.

[65] A. Lytchak and S. Wenger. “Isoperimetric characterization of upper curvature bounds”. *Acta Math.* 221.1 (2018), 159–202.

[66] A. Lytchak. “Differentiation in metric spaces”. *St. Petersburg Math. J.* 16.6 (2005), 1017–1041.

[67] H. von Mangoldt. “Ueber diejenigen Punkte auf positiv gekrümmten Flächen, welche die Eigenschaft haben, dass die von ihnen ausgehenden geodätischen Linien nie aufhören, kürzeste Linien zu sein”. *J. Reine Angew. Math.* 91 (1881), 23–53.

[68] B. Mazur. “A note on some contractible 4-manifolds”. *Ann. of Math. (2)* 73 (1961), 221–228.

[69] K. Menger. „Untersuchungen über allgemeine Metrik“. *Math. Ann.* 100.1 (1928), 75–163.

[70] J. Nagata. *Modern dimension theory.* Revised. Vol. 2. Sigma Series in Pure Mathematics. 1983.

[71] F. Nazarov. *Intrinsic metric with no geodesics.* MathOverflow. eprint: `http://mathoverflow.net/q/15720`.

[72] D. Panov and A. Petrunin. “Sweeping out sectional curvature”. *Geom. Topol.* 18.2 (2014), 617–631.

[73] D. Panov and A. Petrunin. “Ramification conjecture and Hirzebruch’s property of line arrangements”. *Compos. Math.* 152.12 (2016), 2443–2460.

[74] A. Petrunin. *Pure metric geometry.* SpringerBriefs in Mathematics. 2023.

[75] A. Petrunin. “In search of a five-point Aleksandrov type condition”. *St. Petersburg Mathematical Journal* 29.1 (2018), 223–225.

[76] A. Petrunin and S. Stadler. “Metric-minimizing surfaces revisited”. *Geom. Topol.* 23.6 (2019), 3111–3139.

[77] L. Pontrjagin. “Sur une hypothèse fondamentale de la théorie de la dimension.” *C. R. Acad. Sci., Paris* 190 (1930), 1105–1107.

[78] Yu. G. Reshetnyak. “On the theory of spaces with curvature no greater than K”. *Mat. Sb. (N.S.)* 52 (94) (1960), 789–798.

[79] Yu. G. Reshetnyak. “Inextensible mappings in a space of curvature no greater than K”. *Siberian Mathematical Journal* 9.4 (1968), 683–689.

[80] Yu G Reshetnyak. "Two-dimensional manifolds of bounded curvature". *Geometry IV: Non-regular Riemannian Geometry*. Springer, 1993, 3–163.

[81] W. Rinow. *Die innere Geometrie der metrischen Räume*. Die Grundlehren der mathematischen Wissenschaften, Bd. 105. 1961.

[82] D. Rolfsen. "Strongly convex metrics in cells". *Bull. Amer. Math. Soc.* 74 (1968), 171–175.

[83] С. З. Шефель. "О внутренней геометрии седловых поверхностей". *Сибирский математический журнал* 5.6 (1964), 1382–1396.

[84] С. З. Шефель. "О седловых поверхностях, ограниченных спрямляемой кривой". *Доклады Академии наук* 162.2 (1965), 294–296. English translation "On saddle surfaces bounded by a rectifiable curve", Soviet Math. Dokl. 6 (1965), 684–687.

[85] S. Stadler. "An obstruction to the smoothability of singular nonpositively curved metrics on 4-manifolds by patterns of incompressible tori". *Geom. Funct. Anal.* 25.5 (2015), 1575–1587.

[86] D. Stone. "Geodesics in piecewise linear manifolds". *Trans. Amer. Math. Soc.* 215 (1976), 1–44.

[87] P. Thurston. "CAT(0) 4-manifolds possessing a single tame point are Euclidean". *J. Geom. Anal.* 6.3 (1996), 475–494 (1997).

[88] H. Tietze. "Über Konvexheit im kleinen und im großen und über gewisse den Punkten einer Menge zugeordnete Dimensionszahlen". *Math. Z.* 28.1 (1928), 697–707.

[89] S. A. Vakhrameev. "Morse lemmas for smooth functions on manifolds with corners". Vol. 100. 4. Dynamical systems, 8. 2000, 2428–2445.

[90] A. Wald. „Begründung eiiner Koordinatenlosen Differentialgeometrie der Flächen“. *Ergebnisse eines mathematischen Kolloquium* 6 (1935), 24–46.

[91] R. A. Wijsman. "Convergence of sequences of convex sets, cones and functions. II". *Trans. Amer. Math. Soc.* 123.1 (1966), 32–45.