Alexandrov geometry: foundations

Stephanie Alexander, Vitali Kapovitch,
and Anton Petrunin

Preface

Alexandrov spaces are defined via axioms similar to those given by Euclid. The Alexandrov axioms replace certain equalities with inequalities. Depending on the signs of the inequalities, we obtain Alexandrov spaces with curvature bounded above (CBA) and curvature bounded below (CBB). The definitions of the two classes of spaces are similar, but their properties and known applications are quite different. Our approach is novel in its attention to the interrelatedness of the two fields, and its emphasis on the way each illuminates the other.

The goal of this book is to give a comprehensive exposition of the structure theory of Alexandrov spaces with curvature bounded above and below. It includes all the basic material as well as selected topics inspired by considering the two contexts simultaneously. We only consider the intrinsic theory, leaving applications aside.

The book's contents include results that have not appeared elsewhere, as well as many novel approaches.

We note that our presentation is not linear; sometimes proofs and topics are deferred for later chapters to streamline the exposition and make it more natural.

This book includes material up to the definition of dimension. Another volume still in preparation will cover further topics.

Brief history

The first synthetic description of curvature is due to Abraham Wald [128]; it was given in a lone publication on a "coordinateless description of Gauss surfaces" published in 1936. In 1941, similar definitions were rediscovered independently by Aleksandr Alexandrov [10].

In Alexandrov's work the first fruitful applications of this approach were given. Mainly: Alexandrov's embedding theorem, which describes closed convex surfaces in Euclidean 3-space, and the gluing theorem, which gave a flexible tool to modify non-negatively curved metrics on a sphere. These two results together gave a very intuitive geometric tool to study embeddings and bending of surfaces in Euclidean space and changed the subject dramatically. They formed the foundation of the branch of geometry now called Alexandrov geometry.

Curvature bounded below. The theory grew out of studying intrinsic and extrinsic geometry of convex surfaces without the smoothness condition. It was developed by Aleksandr Alexandrov and his school. Here is a very incomplete list of contributors to the subject: Yuriy Borisov, Yuriy Burago, Boris Dekster, Iosif Liberman, Sergey Olovyanishnikov, Aleksey Pogorelov, Yuriy Reshetnyak, Yuriy Volkov, Viktor Zalgaller.

The first result in higher dimensional Alexandrov spaces was the splitting theorem. It was proved by Anatoliy Milka [95] and appeared in 1967. Milka used a global definition similar to the one used in this book.

In the 80's the interest in convergence of Riemannian manifolds spurred by Gromov's compactness theorem [56] turned attention toward the singular spaces that can occur as limits of Riemannian manifolds. Immediately it was recognized that if the manifolds have a uniform lower sectional curvature bound, then the limit spaces have a lower curvature bound in the sense of Alexandrov. There followed during the 90's an explosion of work on intrinsic theory of Alexandrov spaces starting with papers of Yuriy Burago, Grigori Perelman, and Michael Gromov [34, 103]. Similar ideas were developed independently by Karsten Grove and Peter Petersen, whose work was not converted into a publication, and also by Conrad Plaut [110].

Around the same time an implicit application of higher-dimensional Alexandrov geometry was given by Michael Gromov in his bound on Betti numbers [60]. Another implicit application was given later by Wu-Yi Hsiang and Bruce Kleiner in their paper on non-negatively curved manifolds with infinite symmetry groups [68]. The work of Hsiang and Kleiner and its extension by Karsten Grove and Burkhard Wilking [63] are some of the most beautiful applications of this branch of Alexandrov geometry.

The above activity was very much related to so-called comparison geometry, a branch of differential geometry that compares Riemannian manifolds to spaces of constant curvature. In addition to the already-mentioned Gromov's compactness theorem, the following results had a big influence on the development of Alexandrov geometry: Toponogov comparison theorem [125], which is a generalization of the theorem of Alexandrov [9]; Toponogov splitting theorem [125], which is a generalization of Cohn-Vossen's theorem [43]; Finiteness theorems of Cheeger and Grove–Petersen [41, 62]; Gromov's bound on the number of generators of the fundamental group [58, 1.5]; and Yamaguchi fibration theorem [130].

Let us give a list of available introductory texts on Alexandrov spaces with curvature bounded below:

- ⋄ The first introduction to Alexandrov geometry is given in the original paper of Yuriy Burago, Michael Gromov, and Grigori Perelman [34] and its extension [103] written by Perelman.
- ⋄ A brief and reader-friendly introduction was written by Katsuhiro Shiohama [123, Sections 1–8].
- ⋄ [27, Chapter 10] gives another reader-friendly introduction, written by Dmiti Burago, Yuriy Burago, and Sergei Ivanov.

In addition, let us mention two surveys, one by Conrad Plaut [112] and the other by the third author [107].

Curvature bounded above. The study of spaces with curvature bounded above started later, inspired by analogy with the theory of curvature bounded below. The first paper on the subject was written by Alexandrov [12], appearing in 1951. An analogous weaker definition was considered earlier by Herbert Busemann [35].

Contributions to the subject were made by Valerii Berestovskii, Arne Beurling, Igor Nikolaev, Dmitry Sokolov, Yuriy Reshetnyak, Samuel Shefel; this list is not complete as well. The most fundamental results were obtained by Yuriy Reshetnyak. This includes his majorization theorem and gluing theorem. The gluing theorem states that if two nonpositively curved spaces have isometric convex sets, then the space obtained by gluing these sets along an isometry is also non-positively curved.

The development of Alexandrov geometry was greatly influenced by the Hadamard–Cartan theorem. Its original formulation states that the exponential map at any point of a complete Riemannian manifold with nonpositive sectional curvature is a covering. In particular it implies that the universal cover is diffeomorphic to Euclidean space of the same dimension. See further discussion below (9.61).

An influential implicit application of Alexandrov spaces with curvature bounded above can be seen in Euclidean buildings introduced by Jacques Tits as a means to study algebraic groups.

Here is a list of available texts covering the basics of Alexandrov spaces with curvature bounded above:

- ⋄ The book of Martin Bridson and André Haefliger [25] gives the most comprehensive introduction available today.
- ⋄ The lecture notes of Werner Ballmann [15, 16] include a brief and clear introduction.
- ⋄ [27, Chapter 9] gives another reader-friendly introduction, by Yuriy Burago, Dima Burago, and Sergei Ivanov.
- ⋄ A book by the three authors of the present volume [8] gives an introduction aiming at reaching interesting applications and theorems with a minimum of preparation.

- ◇ The book of Jürgen Jost [73] gives a more analytic viewpoint to the subject.

One of the most striking applications of CAT(0) spaces was given by Dmitry Burago, Sergei Ferleger, and Alexey Kononenko [28], who used them to study billiards; this idea was developed further in [29–33]. Another beautiful application is the construction of exotic aspherical manifolds by Michael Davis [46]; related results are surveyed in [39, 47]. Both of these topics are discussed in [8]. The study of group actions on CAT(0) spaces and CAT(0) cube complexes played a key role in the proof of the virtually fibered conjecture that a finite cover of every closed hyperbolic 3-manifold fibers over the circle.

Satellites and successors. Surfaces with bounded integral curvature were studied by Alexandrov's school. An excellent book on the subject was written by Aleksandr Alexandrov and Viktor Zalgaller [1]; see also a more up-to-date survey by Yuriy Reshetnyak [114].

Spaces with two-sided bounded curvature is another subject already studied by Alexandrov's school; a good survey is written by Valerij Berestovskij and Igor Nikolaev [18].

A spin-off of the idea of synthetically defining upper curvature bounds was given by Michael Gromov [61]. He defined so-called δ-hyperbolic spaces, which satisfy a coarse version of the negative curvature condition, applying in particular to discrete metric spaces. This notion and its various generalizations such as semi-hyperbolicity (a coarse version of nonpositive curvature) and relative hyperbolicity have led to the emergence of the subject of geometric group theory which relates geometric properties of groups to their algebraic ones. This is a well-developed subject with a large number of subfields and applications such as the theory of small cancellation groups, automatic groups, mapping class groups, automorphisms of free groups, isoperimetric inequalities on groups, actions on $\mathbb{R}$-trees, Gromov's boundaries of groups.

The so-called curvature dimension condition introduced by John Lott, Cédric Villani, and Karl-Theodor Sturm gives a synthetic description of Ricci curvature bounded below; see the book of Villani [127] and references therein.

Alexandrov geometry influenced the development of analysis on metric spaces. An excellent book on the subject was written by Juha Heinonen, Pekka Koskela, Nageswari Shanmugalingam, and Jeremy Tyson [67].

Acknowledgment

We thank Semyon Alesker, Valerii Berestovskii, I. David Berg, Richard Bishop, Yuriy Burago, Nicola Gigli, Sergei Ivanov, Bernd Kirchheim, Bruce Kleiner, Rostislav Matveyev, John Lott, Alexander Lytchak, Greg Kuperberg, Nikolai Kosovsky, Nina Lebedeva, Wilderich Tuschmann, and Sergio Zamora Barrera.

We thank the mathematical institutions where we worked on this book: BIRS, MFO, Henri Poincaré Institute, University of Cologne, Max Planck Institute for Mathematics.

During the long writing of this book, we were partially supported by the following grants: Stephanie Alexander — Simons Foundation grant 209053; Vitali Kapovitch — NSF grant DMS-0204187, NSERC Discovery grants, and Simons Foundation grant 390117; Anton Petrunin — NSF grants DMS-0406482, DMS-0905138, DMS-1309340, DMS-2005279, and Simons Foundation grants 245094 and 584781.

Contents

Part I

Preliminaries

Chapter 1

Model plane

A Trigonometry

Given a real number κ, the *model κ-plane* will be a complete simply connected 2-dimensional Riemannian manifold of constant curvature κ.

The model κ-plane will be denoted by $\mathbb{M}^2(\kappa)$.

- ◇ If $\kappa > 0$, $\mathbb{M}^2(\kappa)$ is isometric to a sphere of radius $\frac{1}{\sqrt{\kappa}}$; the unit sphere $\mathbb{M}^2(1)$ will be also denoted by $\mathbb{S}^2$.
- ◇ If $\kappa = 0$, $\mathbb{M}^2(\kappa)$ is the Euclidean plane, which is also denoted by $\mathbb{E}^2$.
- ◇ If $\kappa < 0$, $\mathbb{M}^2(\kappa)$ is the Lobachevsky plane with curvature κ.

Set $\varpi^\kappa = \operatorname{diam} \mathbb{M}^2(\kappa)$, so $\varpi^\kappa = \infty$ if $\kappa \leqslant 0$ and $\varpi^\kappa = \pi/\sqrt{\kappa}$ if $\kappa > 0$.

The distance between points $x, y \in \mathbb{M}^2(\kappa)$ will be denoted by $|x - y|$, and $[xy]$ will denote the segment connecting x and y. The segment $[xy]$ is uniquely defined for $\kappa \leqslant 0$ and for $\kappa > 0$ it is defined uniquely if $|x - y| < \varpi^\kappa = \pi/\sqrt{\kappa}$.

A triangle in $\mathbb{M}^2(\kappa)$ with vertices x, y, z will be denoted by $[xyz]$. Formally, a triangle is an ordered set of its sides, so $[xyz]$ is just a short notation for the triple $([yz], [zx], [xy])$.

The angle of $[xyz]$ at x will be denoted by $\measuredangle[x\,{}^y_z]$.

By $\tilde{\triangle}^\kappa\{a, b, c\}$ we denote a triangle in $\mathbb{M}^2(\kappa)$ with side lengths a, b, c, so $[xyz] = \tilde{\triangle}^\kappa\{a, b, c\}$ means that $x, y, z \in \mathbb{M}^2(\kappa)$ are such that

$$|x - y| = c, \quad |y - z| = a, \quad |z - x| = b.$$

For $\tilde{\triangle}^\kappa\{a, b, c\}$ to be defined, the sides a, b, c must satisfy the triangle inequality. If $\kappa > 0$, we require in addition that $a + b + c < 2{\cdot}\varpi^\kappa$; otherwise $\tilde{\triangle}^\kappa\{a, b, c\}$ is considered to be undefined.

Trigonometric functions. We will need three "trigonometric functions" in $\mathbb{M}^2(\kappa)$: cs^κ, sn^κ, and md^κ; cs stands for cosine, sn stands for sine, and md, for modified distance.

They are defined as the solutions of the following initial value problems respectively:

$$\begin{cases} x'' + \kappa{\cdot}x = 0, \\ x(0) = 1, \\ x'(0) = 0. \end{cases} \qquad \begin{cases} y'' + \kappa{\cdot}y = 0, \\ y(0) = 0, \\ y'(0) = 1. \end{cases} \qquad \begin{cases} z'' + \kappa{\cdot}z = 1, \\ z(0) = 0, \\ z'(0) = 0. \end{cases}$$

Namely we set $\operatorname{cs}^\kappa(t) = x(t)$, $\operatorname{sn}^\kappa(t) = y(t)$, and

$$\operatorname{md}^\kappa(t) = \begin{cases} z(t) & \text{if } 0 \leqslant t \leqslant \varpi^\kappa, \\ \frac{2}{\kappa} & \text{if } t > \varpi^\kappa. \end{cases}$$

Here are the tables which relate our trigonometric functions to the standard ones, where we take $\kappa > 0$:

$$\begin{aligned} \operatorname{sn}^{\pm\kappa} &= \tfrac{1}{\sqrt{\kappa}}{\cdot}\operatorname{sn}^{\pm1}(x{\cdot}\sqrt{\kappa}); \\ \operatorname{sn}^{-1}x &= \sinh x; \\ \operatorname{sn}^{0}x &= x; \\ \operatorname{sn}^{1}x &= \sin x; \end{aligned} \qquad \begin{aligned} \operatorname{cs}^{\pm\kappa} &= \operatorname{cs}^{\pm1}(x{\cdot}\sqrt{\kappa}); \\ \operatorname{cs}^{-1}x &= \cosh x; \\ \operatorname{cs}^{0}x &= 1; \\ \operatorname{cs}^{1}x &= \cos x. \end{aligned}$$

$$\begin{aligned} \operatorname{md}^{\pm\kappa} &= \tfrac{1}{\kappa}{\cdot}\operatorname{md}^{\pm1}(x{\cdot}\sqrt{\kappa}); \\ \operatorname{md}^{-1}x &= \cosh x - 1; \\ \operatorname{md}^{0}x &= \tfrac{1}{2}{\cdot}x^2; \\ \operatorname{md}^{1}x &= \left[\begin{array}{lll} 1 - \cos x & \text{for} & x \leqslant \pi, \\ 2 & \text{for} & x > \pi. \end{array}\right. \end{aligned}$$

Note that

$$\operatorname{md}^\kappa(x) = \int_0^x \operatorname{sn}^\kappa(x){\cdot}\boldsymbol{d}x \quad \text{for} \quad x \leqslant \varpi^\kappa$$

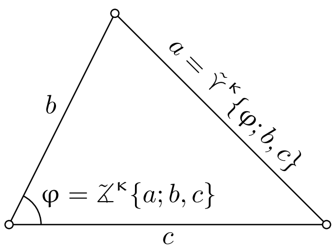

Let φ be the angle of $\tilde{\triangle}^\kappa\{a, b, c\}$ opposite to a. In this case, we will write

$$a = \tilde{\curlyvee}^\kappa\{\varphi; b, c\} \quad \text{or} \quad \varphi = \tilde{\measuredangle}^\kappa\{a; b, c\}.$$

The functions $\tilde{\curlyvee}^\kappa$ and $\tilde{\measuredangle}^\kappa$ will be called respectively the model side and the model angle. Set

$$\tilde{\curlyvee}^\kappa\{\varphi; b, -c\} = \tilde{\curlyvee}^\kappa\{\varphi; -b, c\} := \tilde{\curlyvee}^\kappa\{\pi - \varphi; b, c\};$$

in this way we define $\tilde{\curlyvee}^\kappa\{\varphi; b, c\}$ when one of the numbers b and c is negative.

1.1. Properties of standard functions.

a) *For fixed a and φ, the function*

$$y(t) = \mathrm{md}^\kappa\left(\tilde{\curlyvee}^\kappa\{\varphi; a, t\}\right)$$

satisfies the following differential equation:

$$y'' + \kappa{\cdot}y = 1.$$

b) *Let $\alpha\colon [a, b] \to \mathbb{M}^2(\kappa)$ be a unit-speed geodesic, and A be the image of a complete geodesic. If $f(t)$ is the distance from $\alpha(t)$ to A, the function*

$$y(t) = \mathrm{sn}^\kappa(f(t))$$

satisfies the following differential equation:

$$y'' + \kappa{\cdot}y = 0$$

for $y \neq 0$.

c) *For fixed κ, b, and c, the function*

$$a \mapsto \tilde{\measuredangle}^\kappa\{a; b, c\}$$

is increasing and defined on a real interval. Equivalently, the function

$$\varphi \mapsto \tilde{\curlyvee}^\kappa\{\varphi; b, c\}$$

is increasing and defined if $b, c < \varpi^\kappa$, and $\varphi \in [0, \pi]$[1].

d) *For fixed φ, a, b, c, the function*

$$\kappa \mapsto \tilde{\measuredangle}^\kappa\{a; b, c\} \quad \textit{and} \quad \kappa \mapsto \tilde{\curlyvee}^\kappa\{\varphi; b, c\}$$

are respectively nondereasing (in fact, increasing, if $|b - c| < < a < b + c$) and nonincreasing (in fact, increasing, if $0 < < \varphi < \pi$).

e) *(Alexandrov's lemma) Assume that for real numbers a, b, a', b', x, and κ, the following two expressions are defined:*

---

[1]Formally speaking, if $\kappa > 0$ and $b + c \geqslant \varpi^\kappa$, it is defined only for $\varphi \in [0, \pi)$, but $\tilde{\curlyvee}^\kappa\{\varphi; b, c\}$ can be extended to $[0, \pi]$ as a continuous function.

1. $\measuredangle^\kappa\{a;b,x\} + \measuredangle^\kappa\{a';b',x\} - \pi$,

2. $\measuredangle^\kappa\{a';b+b',a\} - \measuredangle^\kappa\{x;a,b\}$,

Then they have the same sign.

All the properties except Alexandrov's lemma (*e*) can be shown by direct calculation. Alexandrov's lemma is reformulated in 6.2 and is proved there.

Cosine law.

The above formulas easily imply the cosine law in $\mathbb{M}^2(\kappa)$, which can be expressed as follows

$$\cos\varphi = \begin{cases} \dfrac{b^2+c^2-a^2}{2\cdot b\cdot c} & \text{if} \quad \kappa = 0, \\ \dfrac{\operatorname{cs}^\kappa a - \operatorname{cs}^\kappa b\cdot \operatorname{cs}^\kappa c}{\kappa\cdot \operatorname{sn}^\kappa b\cdot \operatorname{sn}^\kappa c} & \text{if} \quad \kappa \neq 0. \end{cases}$$

However, rather than using these explicit formulas, we mainly will use the properties of $\measuredangle^\kappa$ and $\tilde{\curlyvee}^\kappa$ listed in 1.1.

B Hemisphere lemma

1.2. Hemisphere lemma. *For $\kappa > 0$, any closed path of length $< 2\cdot\varpi^\kappa$ (respectively, $\leqslant 2\cdot\varpi^\kappa$) in $\mathbb{M}^2(\kappa)$ lies in an open (respectively, closed) hemisphere.*

Proof. Applying rescaling, we may assume that $\kappa = 1$, and thus $\varpi^\kappa = \pi$ and $\mathbb{M}^2(\kappa) = \mathbb{S}^2$. Let α be a closed curve in $\mathbb{S}^2$ of length $2\cdot\ell$.

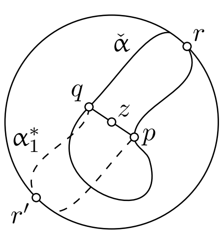

Assume $\ell < \pi$. Let $\breve\alpha$ be a subarc of α of length ℓ, with endpoints p and q. Since $|p-q| \leqslant \leqslant \ell < \pi$, there is a unique geodesic $[pq]$ in $\mathbb{S}^2$. Let z be the midpoint of $[pq]$. We claim that α lies in the open hemisphere centered at z. If not, α intersects the boundary great circle of this hemisphere; let r be a point in the intersection. Without loss of generality we may assume that $r \in \breve\alpha$. The arc $\breve\alpha$ together with its reflection in z form a closed curve of length $2\cdot\ell$ that contains r and its antipodal point r'. Thus

$$\ell = \operatorname{length}\breve\alpha \geqslant |r-r'| = \pi,$$

a contradiction.

If $\ell = \pi$, then either α is a local geodesic, and hence a great circle, or α may be strictly shortened by substituting a geodesic arc for a subarc of α whose endpoints p^1, p^2 are arbitrarily close to some point p on α. In the latter case, α lies in a closed hemisphere obtained as a limit of closures of open hemispheres containing the shortened curves as p^1, p^2 approach p. □

1.3. Exercise. *Give a proof of the hemisphere lemma (1.2) based on Crofton's formula.*

Chapter 2

Metric spaces

In this chapter we fix some conventions and notations. We are assuming that the reader is familiar with basic notions in metric geometry.

A Metrics and their relatives

Definitions. Let $\mathbb{I}$ be a subinterval of $[0,\infty]$. A function ρ defined on $\mathcal{X}\times\mathcal{X}$ is called an $\mathbb{I}$-valued metric if the following conditions hold:

- ◇ $\rho(x,x)=0$ for any x;
- ◇ $\rho(x,y)\in\mathbb{I}$ for any pair $x\neq y$;
- ◇ $\rho(x,y)+\rho(x,z)\geqslant\rho(y,z)$ for any triple of points x,y,z.

The value $\rho(x,y)$ is also called the distance between x and y.

The above definition will be used for four choices of interval $\mathbb{I}$: $(0,\infty)$, $(0,\infty]$, $[0,\infty)$, and $[0,\infty]$. Any $\mathbb{I}$-valued metric can be referred to briefly as a metric; the interval should be apparent from context but by default, a metric is $(0,\infty)$-valued. If we need to be more specific we may also use the following names:

- ◇ a $(0,\infty)$-valued metric may be called a genuine metric.
- ◇ a $(0,\infty]$-valued metric may be called an ∞-metric.
- ◇ a $[0,\infty)$-valued metric may be called a genuine pseudometric.
- ◇ A $[0,\infty]$-valued metric may be called a pseudometric or ∞-pseudometric.

A metric space is a set equipped with a metric. The distance between points x and y in a metric space $\mathcal{X}$ will usually be denoted by

$$|x-y| \qquad \text{or} \quad |x-y|_{\mathcal{X}};$$

the latter will be used if we need to emphasize that we are working in the space $\mathcal{X}$.

The function $\operatorname{dist}_x : \mathcal{X} \to \mathbb{R}$ defined as

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

will be called the distance function from x.

Any subset A in a metric space $\mathcal{X}$ will be also considered as a subspace; that is, a metric space with the metric defined by restricting the metric of $\mathcal{X}$ to $A \times A \subset \mathcal{X} \times \mathcal{X}$.

The direct product $\mathcal{X} \times \mathcal{Y}$ of two metric spaces $\mathcal{X}$ and $\mathcal{Y}$ is defined as the metric space carrying the metric

$$|(p, \varphi) - (q, \psi)| = \sqrt{|p - q|^2 + |\varphi - \psi|^2}$$

for $p, q \in \mathcal{X}$ and $\varphi, \psi \in \mathcal{Y}$.

Zero and infinity. Genuine metric spaces are the main objects of study in this book. However, the generalizations above are useful in various definitions and constructions. For example, the construction of length metric (see Section 2C) uses infinite distances. The following definition gives another example.

2.1. Definition. *Assume $\{\mathcal{X}_\alpha\}_{\alpha \in \mathcal{A}}$ is a collection of ∞-metric spaces. The disjoint union*

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

has a natural ∞-metric on it defined as follows: given two points $x \in \mathcal{X}_\alpha$ and $y \in \mathcal{X}_\beta$ set

$$\begin{aligned} |x - y|_{\boldsymbol{X}} &= \infty & &\text{if} \quad \alpha \neq \beta, \\ |x - y|_{\boldsymbol{X}} &= |x - y|_{\mathcal{X}_\alpha} & &\text{if} \quad \alpha = \beta. \end{aligned}$$

The resulting ∞-metric space $\boldsymbol{X}$ will be called the disjoint union of $\{\mathcal{X}_\alpha\}_{\alpha \in \mathcal{A}}$, denoted by

$$\bigsqcup_{\alpha \in \mathcal{A}} \mathcal{X}_\alpha.$$

Now let us give examples showing that vanishing and infinite distance between distinct points is useful.

Suppose a set $\mathcal{X}$ comes with a set of metrics $|{*} - {*}|_\alpha$ for $\alpha \in \mathcal{A}$. Then

$$|x - y| = \sup\{\, |x - y|_\alpha : \alpha \in \mathcal{A} \,\}$$

is in general only an ∞-metric; that is, even if the metrics $|{*} - {*}|_\alpha$ are genuine, then $|{*} - {*}|$ might be $(0, \infty]$-valued.

Let $\mathcal{X}$ be a set, $\mathcal{Y}$ be a metric space, and $\Phi\colon \mathcal{X} \to \mathcal{Y}$ be a map. If Φ is not injective, then the pullback

$$|x - y|_{\mathcal{X}} = |\Phi(x) - \Phi(y)|_{\mathcal{Y}}$$

defines only a pseudometric on $\mathcal{X}$.

Corresponding metric space and metric component. The following two observations show that nearly any question about metric spaces can be reduced to a question about genuine metric spaces.

Assume $\mathcal{X}$ is a pseudometric space. Set $x \sim y$ if $|x - y| = 0$. Note that if $x \sim x'$, then $|y - x| = |y - x'|$ for any $y \in \mathcal{X}$. Thus, $|{*} - {*}|$ defines a metric on the quotient set $\mathcal{X}/\sim$. In this way we obtain a metric space $\mathcal{X}'$. The space $\mathcal{X}'$ is called the corresponding metric space for the pseudometric space $\mathcal{X}$. Often we do not distinguish between $\mathcal{X}'$ and $\mathcal{X}$.

Set $x \approx y$ if and only if $|x-y| < \infty$; this is another equivalence relation on $\mathcal{X}$. The equivalence class of a point $x \in \mathcal{X}$ will be called the metric component of x; it will be denoted as $\mathcal{X}_x$. One could think of $\mathcal{X}_x$ as $\mathrm{B}(x, \infty)_{\mathcal{X}}$, the open ball centered at x and radius ∞ in $\mathcal{X}$; see definition below.

It follows that any ∞-metric space is a disjoint union of genuine metric spaces, the metric components of the original ∞-metric space; see Definition 2.1

To summarize this discussion: Given a $[0, \infty]$-valued metric space $\mathcal{X}$, we may pass to the corresponding $(0, \infty]$-valued metric space $\mathcal{X}'$ and break the latter into a disjoint union of metric components, each of which is a genuine metric space.

B Notations

Balls. Given $R \in [0, \infty]$ and a point x in a metric space $\mathcal{X}$, the sets

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

are called respectively the open and the closed balls of radius R with center x.

If we need to emphasize that these balls are taken in the space $\mathcal{X}$, we write

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

respectively.

Since in the model space $\mathbb{M}^m(\kappa)$ all balls of the same radius are isometric, often we will not need to specify the center of the ball, and may write

$$\mathrm{B}(R)_{\mathbb{M}^m(\kappa)} \quad \text{and} \quad \overline{\mathrm{B}}[R]_{\mathbb{M}^m(\kappa)}$$

respectively.

A set $A \subset \mathcal{X}$ is called *bounded* if $A \subset \mathrm{B}(x, R)$ for some $x \in \mathcal{X}$ and $R < \infty$.

Distances to sets. For subset $A \subset \mathcal{X}$, let us denote the distance from A to a point x in $\mathcal{X}$ as $\operatorname{dist}_A x$; that is,

$$\operatorname{dist}_A x := \inf \{ |a - x| : a \in A \} .$$

We define the distance between sets A and B as

$$|A - B| := \inf \{ |a - b| : a \in A, \ b \in B \} .$$

For any subset $A \subset \mathcal{X}$, the sets

$$\begin{aligned} \mathrm{B}(A, R) &= \{ y \in \mathcal{X} : \operatorname{dist}_A y < R \} , \\ \overline{\mathrm{B}}[A, R] &= \{ y \in \mathcal{X} : \operatorname{dist}_A y \leqslant R \} \end{aligned}$$

are called respectively the *open* and *closed* R-*neighborhoods* of A.

Diameter, radius, and packing. Let $\mathcal{X}$ be a metric space. Then the *diameter* of $\mathcal{X}$ is defined as

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

The *radius* of $\mathcal{X}$ is defined as

$$\operatorname{rad} \mathcal{X} = \inf \{ R > 0 : \mathrm{B}(x, R) = \mathcal{X} \text{ for some } x \in \mathcal{X} \} .$$

The packing number ε-*pack* of $\mathcal{X}$ is the maximal number (possibly infinite) of points in $\mathcal{X}$ at distance $> \varepsilon$ from each other; it is denoted by $\operatorname{pack}_\varepsilon \mathcal{X}$. If $m = \operatorname{pack}_\varepsilon \mathcal{X} < \infty$, then a set $\{x^1, x^2, \dots, x^m\}$ in $\mathcal{X}$ such that $|x^i - x^j| > \varepsilon$ is called a *maximal* ε-*packing* in $\mathcal{X}$.

G-delta sets. Recall that an arbitrary union of open balls in a metric space is called an *open set*. A subset of a metric space is called a *G-delta set* if it can be presented as an intersection of a countable number of open subsets.

Often we will use the following classical result:

2.2. Baire's theorem. *Let $\mathcal{X}$ be a complete metric space and $\{\Omega_n\}$, $n \in \mathbb{N}$, be a collection of open dense subsets of $\mathcal{X}$. Then $\bigcap_{n=1}^{\infty} \Omega_n$ is dense in $\mathcal{X}$.*

Proper spaces. 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 $\mathrm{dist}_p : \mathcal{X} \to \mathbb{R}$ is proper for some (and therefore any) point $p \in \mathcal{X}$.

We will often use the following two classical statements:

2.3. Proposition. *Proper metric spaces are separable and second countable.*

2.4. Proposition. *Let $\mathcal{X}$ be a metric space. Then the following are equivalent*

i) $\mathcal{X}$ is compact;

ii) $\mathcal{X}$ is sequentially compact; that is, any sequence of points in $\mathcal{X}$ contains a convergent subsequence;

iii) $\mathcal{X}$ is complete and for any $\varepsilon > 0$ there is a finite ε-net in $\mathcal{X}$; that is, there is a finite collection of points $p_1, \dots, p_N$ such that $\bigcup_i \mathrm{B}(p_i, \varepsilon) = \mathcal{X}$.

C Length spaces

A curve in $\mathcal{X}$ is a continuous map $\alpha \colon \mathbb{I} \to \mathcal{X}$, where $\mathbb{I}$ is a real interval (that is, an arbitrary convex subset of $\mathbb{R}$).

2.5. Definition. *Let $\mathcal{X}$ be a metric space. Given a curve $\alpha \colon \mathbb{I} \to \mathcal{X}$, we define its length as*

$$\operatorname{length}\alpha := \sup\left\{\sum_{i \geqslant 1} |\alpha(t_i) - \alpha(t_{i-1})| \ : \ t_0, \dots, t_n \in \mathbb{I}, t_0 \leqslant \dots \leqslant t_n\right\}$$

The following lemma is an easy exercise.

2.6. Lower semicontinuity of length. *Assume $\alpha_n \colon \mathbb{I} \to \mathcal{X}$ is a sequence of curves that converges pointwise to a curve $\alpha_\infty \colon \mathbb{I} \to \mathcal{X}$. Then*

$$\operatorname{length}\alpha_\infty \leqslant \varliminf_{n \to \infty} \operatorname{length}\alpha_n.$$

Given two points x and y in a metric space $\mathcal{X}$, consider the value

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

where infimum is taken for all paths α from x to y.

It is easy to see that $\|*-*\|$ defines a $(0,\infty]$-valued metric on $\mathcal{X}$; it will be called the induced length metric on $\mathcal{X}$. Clearly

$$\|x-y\| \geqslant |x-y|$$

for any $x, y \in \mathcal{X}$.

It easily follows from the definition that the length of a curve α with respect to $\|*-*\|$ is equal to the length of α with respect to $|*-*|$. In particular, iterating the construction produces the same $\|*-*\|$.

2.7. Definition. *If $\|x-y\| = |x-y|$ for any pair of points x, y in a metric space $\mathcal{X}$, then $\mathcal{X}$ is called a length space.*

In other words, a metric space $\mathcal{X}$ is a length space if for any $\varepsilon > 0$ and any two points $x, y \in \mathcal{X}$ with $|x-y| < \infty$ there is a path $\alpha\colon [0,1] \to \mathcal{X}$ connecting[1] x to y such that

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

In this book, most of the time we consider length spaces. If $\mathcal{X}$ is a length space, and $A \subset \mathcal{X}$ the set A comes with the inherited metric from $\mathcal{X}$ which might be not a length metric. The corresponding length metric on A will be denoted as $|*-*|_A$.

Variations of the definition. We will need the following variations of Definition 2.7:

- ◇ Assume $R > 0$. If $\|x-y\| = |x-y|$ for any pair $|x-y| < R$, then $\mathcal{X}$ is called an R-length space.
- ◇ If any point in $\mathcal{X}$ admits a neighborhood Ω such that $\|x-y\| = |x-y|$ for any pair of points $x, y \in \Omega$ then $\mathcal{X}$ is called a locally length space.
- ◇ A metric space $\mathcal{X}$ is called geodesic if for any two points $x, y \in \mathcal{X}$ with $|x-y| < \infty$ there is a geodesic $[xy]$ in $\mathcal{X}$.
- ◇ Assume $R > 0$. A metric space $\mathcal{X}$ is called R-geodesic if for any two points $x, y \in \mathcal{X}$ such that $|x-y| < R$ there is a geodesic $[xy]$ in $\mathcal{X}$.

Note that the notions of ∞-length spaces and length spaces are the same. Clearly, any geodesic space is a length space and any R-geodesic space is R-length.

2.8. 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 $\mathbb{I}_n$ one end is glued to p and the other to q. Then $\mathcal{X}$ carries a natural complete length metric such that $|p-q| = 1$, but there is no geodesic connecting p to q.*

2.9. Exercise. *Let $\mathcal{X}$ be a metric space and $\|*-*\|$ be the length metric on it. Show the following:*

[1] That is, such that $\alpha(0) = x$ and $\alpha(1) = y$.

a) If $\mathcal{X}$ is complete, then $(\mathcal{X}, \|{*}-{*}\|)$ is complete.
b) If $\mathcal{X}$ is compact, then $(\mathcal{X}, \|{*}-{*}\|)$ is geodesic.

2.10. Exercise. *Give an example of a complete length space such that no pair of distinct points can be joined by a geodesic.*

2.11. Definition. *Let $\mathcal{X}$ be a metric space and $x, y \in \mathcal{X}$.*
(i) A point $z \in \mathcal{X}$ is called a m i d p o i n t *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 - m i d p o i n t *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.

2.12. 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.
c) If for some $R > 0$, the assumptions (a) or (b) hold only for pairs of points $x, y \in \mathcal{X}$ such that $|x - y| < R$, then $\mathcal{X}$ is an R-length or R-geodesic space respectively.

Proof. Fix a pair of points $x, y \in \mathcal{X}$. Let $\varepsilon_n = \frac{\varepsilon}{2^{2\cdot n}}$, $\alpha(0) = x$, 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 $0 < \frac{k}{2^n} < 1$, to be 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 all dyadic rationals t in $[0,1]$. If $t \in [0,1]$ is not a dyadic rational, consider a sequence of dyadic rationals $t_n \to t$ as $n \to \infty$. By completeness of $\mathcal{X}$, the sequence $\alpha(t_n)$ converges; let $\alpha(t)$ be its limit. It is easy to see that $\alpha(t)$ does not depend on the choice of the sequence t_n, and $\alpha\colon [0,1] \to \mathcal{X}$ is a path from x to y. 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 have (a).

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

The proof of (2.12c) is obtained by a straightforward modification of the proofs above. □

2.13. Corollary. *A proper length space is geodesic.*

This follows from Lemma 2.12 since in a compact set a sequence of $\frac{1}{n}$-midpoints (z_n) contains a convergent subsequence.

2.14. 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, $\rho(x) < \infty$.

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 is compact by Proposition 2.4. △

❹ $|\rho(x) - \rho(y)| \leqslant |x - y|_{\mathcal{X}}$, *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 $\varepsilon = \min_{y \in B}\{\rho(y)\}$; the minimum is defined since B is compact. From ❷, we have $\varepsilon > 0$.

Choose a finite $\frac{\varepsilon}{10}$-net $\{a_1, a_2, \ldots, a_n\}$ 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. □

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

D Convex sets

2.16. Definition. *Let $\mathcal{X}$ be a geodesic space and $A \subset \mathcal{X}$.*

We say that A is convex if for every two points $p, q \in A$ any geodesic $[pq]$ lies in A.

We say that A is weakly convex if for every two points $p, q \in A$ there is a geodesic $[pq]$ that lies in A.

We say A is totally convex if for every two points $p, q \in A$, every local geodesic from p to q lies in A.

If for some $R \in (0, \infty]$ these definitions are applied only for pairs of points such that $|p - q| < R$ and only for the geodesics of length $< R$, then A is called respectively R-convex, weakly R-convex, or totally R-convex.

A set $A \subset \mathcal{X}$ is called locally convex if every point $a \in A$ admits an open neighborhood $\Omega \ni a$ such that for every two points $p, q \in A \cap \Omega$ every geodesic $[pq] \subset \Omega$ lies in A. Similarly one defines locally weakly convex and locally totally convex sets.

Remarks. Let us state a few observations that easily follow from the definition.

- ◇ The notion of (weakly) convex set is the same as (weakly) ∞-convex set.
- ◇ The inherited metric on a weakly convex set coincides with its length metric.
- ◇ Any open set is locally convex by definition.

The following proposition states that weak convexity survives under ulralimits. An analogous statement about convexity does not hold; for example there is a sequence of convex discs in $\mathbb{S}^2$ that converges to the hemisphere $\mathbb{S}^2_+$, which is not convex.

2.17. Proposition. *Let $\mathcal{X}_n$, be a sequence of geodesic spaces. Let ω be an ultrafilter on $\mathbb{N}$; see Definition 3.1. Assume that $A_n \subset \mathcal{X}_n$ is a sequence of weakly convex sets, $\mathcal{X}_n \to \mathcal{X}_\omega$, and $A_n \to A_\omega \subset \mathcal{X}_\omega$ as $n \to \omega$. Then A_ω is a weakly convex set of $\mathcal{X}_\omega$.*

Proof. Fix $x_\omega, y_\omega \in A_\omega$. Consider sequences $x_n, y_n \in A_n$ such that $x_n \to \to x_\omega$ and $y_n \to y_\omega$ as $n \to \omega$.

Denote by α_n a geodesic path from x_n to y_n that lies in A_n. Set

$$\alpha_\omega(t) = \lim_{n \to \omega} \alpha_n(t).$$

Note that α_ω is a geodesic path that lies in A_ω. The proposition follows. □

E Quotient spaces

Quotient spaces. Assume $\mathcal{X}$ is a metric space with an equivalence relation $\sim$. Note that given two pseudometrics ρ_1 and ρ_2 on $\mathcal{X}/\sim$, their maximum

$$\rho(x,y) = \max\{\rho_1(x,y), \rho_2(x,y)\}$$

is also a pseudometric. If for these two pseudometrics ρ_1 and ρ_2 the projections $\mathcal{X} \to (\mathcal{X}/\sim, \rho_i)$ are short (that is, distance non-increasing), then the same is true for $\rho(x,y)$.

It follows that the quotient space $\mathcal{X}/\sim$ admits a natural quotient pseudometric; this is the maximal pseudometric on $\mathcal{X}/\sim$ that makes the quotient map $\mathcal{X} \to \mathcal{X}/\sim$ short. The corresponding metric space will be also denoted as $\mathcal{X}/\sim$ and will be called the quotient space of $\mathcal{X}$ by the equivalence relation $\sim$.

In general, the points of the metric space $\mathcal{X}/\sim$ are formed by equivalence classes in $\mathcal{X}$ for a wider equivalence relation. However, in most of the cases we will consider, the set of equivalence classes will coincide with the set of points in the metric space $\mathcal{X}/\sim$.

2.18. Proposition. *Let $\mathcal{X}$ be a length space and $\sim$ be an equivalence relation on $\mathcal{X}$. Then $\mathcal{X}/\sim$ is a length space.*

Proof. Let $\mathcal{Y}$ be an arbitrary metric space. Since $\mathcal{X}$ is a length space, a map $f\colon \mathcal{X} \to \mathcal{Y}$ is short if and only if

$$\operatorname{length}(f \circ \alpha) \leqslant \operatorname{length} \alpha$$

for any curve $\alpha\colon \mathbb{I} \to \mathcal{X}$. Denote by $\|{*} - {*}\|$ the length metric on $\mathcal{Y}$. It follows that if $f\colon \mathcal{X} \to \mathcal{Y}$ is short then so is $f\colon \mathcal{X} \to (\mathcal{Y}, \|{*} - {*}\|)$.

Consider the quotient map $f\colon \mathcal{X} \to \mathcal{X}/\sim$. Recall that the space $\mathcal{X}/\sim$ is defined by the maximal pseudometric that makes f short.

Denoting by $\|{*} - {*}\|$ the length metric on $\mathcal{X}/\sim$, it follows that

$$f\colon \mathcal{X} \to (\mathcal{X}/\sim, \|{*} - {*}\|)$$

is also short.

Note that

$$\|x - y\| \geqslant |x - y|_{\mathcal{X}/\sim}$$

for any $x, y \in \mathcal{X}/\sim$. From maximality of $|{*} - {*}|_{\mathcal{X}/\sim}$, we get

$$\|x - y\| = |x - y|_{\mathcal{X}/\sim}$$

for any $x, y \in \mathcal{X}/\sim$; that is, $\mathcal{X}/\sim$ is a length space. □

Group actions. Assume a group G acts on a metric space $\mathcal{X}$. Consider a relation $\sim$ on $\mathcal{X}$ defined by $x \sim y$ if there is $g \in G$ such that $x = g{\cdot}y$. Note that $\sim$ is an equivalence relation.

In this case, the quotient space $\mathcal{X}/{\sim}$ will also be denoted by $\mathcal{X}/G$, and can be regarded as the space of G-orbits in $\mathcal{X}$.

Assume that a group G acts on $\mathcal{X}$ by isometries. Then the distance between orbits $G{\cdot}x$ and $G{\cdot}y$ in $\mathcal{X}/G$ can be defined directly:

$$|G{\cdot}x - G{\cdot}y|_{\mathcal{X}/G} = \inf\left\{\, |x - g{\cdot}y|_{\mathcal{X}} = |g^{-1}{\cdot}x - y|_{\mathcal{X}} \,:\, g \in G \,\right\}.$$

If the G-orbits are closed, then $|G{\cdot}x - G{\cdot}y|_{\mathcal{X}/G} = 0$ if and only if $G{\cdot}x = G{\cdot}y$. In this case, the quotient space $\mathcal{X}/G$ is a genuine metric space.

The following proposition follows from the definition of a quotient space:

2.19. Proposition. *Assume $\mathcal{X}$ is a metric space and a group G acts on $\mathcal{X}$ by isometries. Then the projection $\pi\colon \mathcal{X} \to \mathcal{X}/G$ is a submetry; that is, $\pi(\mathrm{B}(p,r)) = \mathrm{B}(\pi(p),r)$ for any $p \in \mathcal{X}, r > 0$ (see Definition 5.7).*

F Gluing and doubling

Gluing. Recall that the disjoint union of metric spaces can be also considered as a metric space; see Definition 2.1. Therefore the quotient space construction works as well for an equivalence relation on the disjoint union of metric spaces.

Consider two metric spaces $\mathcal{X}_1$ and $\mathcal{X}_2$ with subsets $A_1 \subset \mathcal{X}_1$ and $A_2 \subset \mathcal{X}_2$, and a bijection $\varphi\colon A_1 \to A_2$. Consider the minimal equivalence relation on $\mathcal{X}_1 \sqcup \mathcal{X}_2$ such that $a \sim \varphi(a)$ for any $a \in A_1$. In this case, the corresponding quotient space $(\mathcal{X}_1 \sqcup \mathcal{X}_2)/{\sim}$ will be called the *gluing of $\mathcal{X}$ and $\mathcal{Y}$ along φ* and denoted by

$$\mathcal{X}_1 \sqcup_\varphi \mathcal{X}_2.$$

Note that if the map $\varphi\colon A_1 \to A_2$ is distance-preserving, then the projections $\iota_i\colon \mathcal{X}_i \to \mathcal{X}_1 \sqcup_\varphi \mathcal{X}_2$ are also distance-preserving, and

$$|\iota_1(x_1) - \iota_2(x_2)|_{\mathcal{X}_1 \sqcup_\varphi \mathcal{X}_2} = \inf_{a_2 = \varphi(a_1)} \left\{\, |x_1 - a_1|_{\mathcal{X}_1} + |x_2 - a_2|_{\mathcal{X}_2} \,\right\}$$

for any $x_1 \in \mathcal{X}_1$ and $x_2 \in \mathcal{X}_2$.

Doubling. Let $\mathcal{V}$ be a metric space and $A \subset \mathcal{V}$ be a closed subset. A metric space $\mathcal{W}$ glued from two copies of $\mathcal{V}$ along A is called the *doubling of $\mathcal{V}$ in A*.

The space $\mathcal{W}$ is completely described by the following properties:

- ⋄ The space $\mathcal{W}$ contains $\mathcal{V}$ as a subspace; in particular the set A can be treated as a subset of $\mathcal{W}$.
- ⋄ There is an isometric involution of $\mathcal{W}$ which is called reflection in A; further it will be denoted as $x \mapsto x'$.
- ⋄ For any $x \in \mathcal{W}$ we have $x \in \mathcal{V}$ or $x' \in \mathcal{V}$ and

$$|x' - y|_{\mathcal{W}} = |x - y'|_{\mathcal{W}} = \inf_{a \in A} \{|x - a|_{\mathcal{V}} + |a - y|_{\mathcal{V}}\}$$

for any $x, y \in \mathcal{V}$.

The image of $\mathcal{V}$ under reflection in A will be denoted by $\mathcal{V}'$. The subspace $\mathcal{V}'$ is an isometric copy of $\mathcal{V}$. Clearly $\mathcal{V} \cup \mathcal{V}' = \mathcal{W}$ and $\mathcal{V} \cap \cap \mathcal{V}' = A$. Moreover $a = a' \iff a \in A$.

The following proposition follows directly from the definitions.

2.20. Proposition. *Assume $\mathcal{W}$ is the doubling of the metric space $\mathcal{V}$ in its closed subset A. Then:*

a) If $\mathcal{V}$ is a complete length space, then so is $\mathcal{W}$.

b) If $\mathcal{V}$ is proper, then so is $\mathcal{W}$. In this case, for any $x, y \in \mathcal{V}$ there is $a \in A$ such that

$$|x - a|_{\mathcal{V}} + |a - y|_{\mathcal{V}} = |x - y'|_{\mathcal{W}}.$$

c) Given $x \in \mathcal{W}$, set $\bar{x} = x$ if $x \in \mathcal{V}$, and $\bar{x} = x'$ otherwise. The map $\mathcal{W} \to \mathcal{V}$ defined by $x \mapsto \bar{x}$ is short and length-preserving. In particular, if γ is a geodesic in $\mathcal{W}$ with ends in $\mathcal{V}$, then $\bar{\gamma}$ is a geodesic in $\mathcal{V}$ with the same ends.

G Kuratowsky embedding

Given a metric space $\mathcal{X}$, let us denote by $\mathrm{Bnd}(\mathcal{X}, \mathbb{R})$ the space of all bounded functions on $\mathcal{X}$ equipped with the sup-norm

$$\|f\| = \sup_{x \in \mathcal{X}} \{|f(x)|\}.$$

Kuratowski embedding. Given a point $p \in \mathcal{X}$, consider the map $\mathrm{kur}_p \colon \mathcal{X} \to \mathrm{Bnd}(\mathcal{X}, \mathbb{R})$ defined by $\mathrm{kur}_p x = \mathrm{dist}_x - \mathrm{dist}_p$. The map kur_p will be called the Kuratowski map at p.

From the triangle inequality, we have

$$\|\mathrm{kur}_p x - \mathrm{kur}_p y\| = \sup_{z \in \mathcal{X}} \{||x - z| - |y - z||\} = |x - y|.$$

Therefore, for any $p \in \mathcal{X}$, the Kuratowski map gives a distance-preserving map $\mathrm{kur}_p \colon \mathcal{X} \hookrightarrow \mathrm{Bnd}(\mathcal{X}, \mathbb{R})$. Thus we can (and often will) consider the space $\mathcal{X}$ as a subset of $\mathrm{Bnd}(\mathcal{X}, \mathbb{R})$.

2.21. Exercise. *Show that any compact metric space is isometric to a subspace in a compact length space.*

Chapter 3

Ultralimits

Here we introduce ultralimits of sequences of points, metric spaces, and functions. The ultralimits of metric spaces can be considered as a variation of Gromov–Hausdorff convergence. Our presentation is based on [78].

Our use of ultralimits is very limited; we use them only as a canonical way to pass to convergent subsequences. (They could be avoided, but at the cost of saying "pass to a convergent subsequence" too many times.)

A Ultrafilters

We will need the existence of a fixed selective ultrafilter ω. The existence follows from the axiom of choice and continuum hypothesis.

Recall that $\mathbb{N}$ denotes the set of natural numbers, $\mathbb{N} = \{1, 2, \ldots\}$.

3.1. Definition. *A finitely additive measure ω on $\mathbb{N}$ is called an ultrafilter if it satisfies*

a) $\omega(S) = 0$ or 1 for any subset $S \subset \mathbb{N}$.

An ultrafilter ω is called nonprincipal if in addition

b) $\omega(F) = 0$ for any finite subset $F \subset \mathbb{N}$.

A nonprincipal ultrafilter ω is called selective if in addition

c) for any partition of $\mathbb{N}$ into sets $\{C_\alpha\}_{\alpha\in\mathcal{A}}$ such that $\omega(C_\alpha) = 0$ for each α, there is a set $S \subset \mathbb{N}$ such that $\omega(S) = 1$ and $S \cap C_\alpha$ is a one-point set for each $\alpha \in \mathcal{A}$.

If $\omega(S) = 0$ for some subset $S \subset \mathbb{N}$, we say that S is ω-small. If $\omega(S) = 1$, we say that S contains ω-almost all elements of $\mathbb{N}$.

Classical definition. More commonly, a nonprincipal ultrafilter is defined as a collection, say $\mathfrak{F}$, of sets in $\mathbb{N}$ such that

1. if $P \in \mathfrak{F}$ and $Q \supset P$, then $Q \in \mathfrak{F}$,
2. if $P, Q \in \mathfrak{F}$, then $P \cap Q \in \mathfrak{F}$,
3. for any subset $P \subset \mathbb{N}$, either P or its complement is an element of $\mathfrak{F}$,
4. if $F \subset \mathbb{N}$ is finite, then $F \notin \mathfrak{F}$.

Setting

$$P \in \mathfrak{F} \quad \Longleftrightarrow \quad \omega(P) = 1$$

makes these two definitions equivalent.

A nonempty collection of sets $\mathfrak{F}$ that does not include the empty set and satisfies only conditions 1 and 2 is called a filter; if in addition $\mathfrak{F}$ satisfies Condition 3 it is called an ultrafilter. From Zorn's lemma, it follows that every filter contains an ultrafilter. Thus there is an ultrafilter $\mathfrak{F}$ contained in the filter of all complements of finite sets; clearly this $\mathfrak{F}$ is nonprincipal.

The existence of a selective ultrafilter follows from the continuum hypothesis; this was proved by Walter Rudin [117].

Stone–Čech compactification. Given a set $S \subset \mathbb{N}$, consider the subset Ω_S of all ultrafilters ω such that $\omega(S) = 1$. It is straightforward to check that the sets Ω_S for all $S \subset \mathbb{N}$ form a topology on the set of ultrafilters on $\mathbb{N}$. The resulting space is called the Stone–Čech compactification of $\mathbb{N}$; it is usually denoted by $\beta\mathbb{N}$.

There is a natural embedding $\mathbb{N} \hookrightarrow \beta\mathbb{N}$ defined by $n \mapsto \omega_n$, where ω_n is the principal ultrafilter such that $\omega_n(S) = 1$ if and only if $n \in S$. Using this embedding, we can (and will) consider $\mathbb{N}$ as a subset of $\beta\mathbb{N}$.

The space $\beta\mathbb{N}$ is the maximal compact Hausdorff space that contains $\mathbb{N}$ as an everywhere dense subset. More precisely, for any compact Hausdorff space $\mathcal{X}$ and map $f\colon \mathbb{N} \to \mathcal{X}$, there is a unique continuous map $\bar{f}\colon \beta\mathbb{N} \to$ $\to X$ such that the restriction $\bar{f}|_{\mathbb{N}}$ coincides with f.

B Ultralimits of points

Fix an ultrafilter ω. Assume x_n is a sequence of points in a metric space $\mathcal{X}$. Define an ω-limit of x_n to be a point x_ω such that for any $\varepsilon > 0$, ω-almost all elements of x_n lie in $\mathrm{B}(x_\omega, \varepsilon)$; that is,

$$\omega\{\, n \in \mathbb{N} : |x_\omega - x_n| < \varepsilon \,\} = 1.$$

In this case, we write

$$x_\omega = \lim_{n \to \omega} x_n \quad \text{or} \quad x_n \to x_\omega \quad \text{as} \quad n \to \omega.$$

Also, if $\mathcal{X} = \mathbb{R}$ we write

$$\lim_{n \to \omega} x_n = \pm\infty$$

if

$$\omega\{\, n \in \mathbb{N} : \pm|x_\omega - x_n| > L \,\} = 1 \quad \text{for any} \quad L \in \mathbb{R}.$$

It easily follows from the definition that ω-limits are unique if they exist. For example if ω is the principal ultrafilter such that $\omega(\{n\}) = 1$ for some $n \in \mathbb{N}$, then $x_\omega = x_n$.

Note that ω-limits of a sequence and its subsequences may differ. For example, in general

$$\lim_{n\to\omega} x_n \neq \lim_{n\to\omega} x_{2\cdot n}.$$

The sequence x_n can be regarded as a map $\mathbb{N} \to \mathcal{X}$. If $\mathcal{X}$ is compact, then this map can be uniquely extended to a continuous map to the Stone–Čech compactification $\beta\mathbb{N}$ of $\mathbb{N}$. Then x_ω can be regarded as the image of ω.

3.2. Proposition. *Let ω be a nonprincipal ultrafilter. Assume x_n is a sequence of points in a metric space $\mathcal{X}$ and $x_n \to x_\omega$ as $n \to \omega$. Then there is a Cauchy subsequence of x_n that converges to x_ω in the usual sense.*

Moreover, if ω is selective, then the subsequence $(x_n)_{n\in S}$ can be chosen so that $\omega(S) = 1$.

Proof. Given $\varepsilon > 0$, let $S_\varepsilon = \{\, n \in \mathbb{N} : |x_n - x_\omega| < \varepsilon \,\}$.

Note that $\omega(S_\varepsilon) = 1$ for any $\varepsilon > 0$. Since ω is nonprincipal, the set S_ε is infinite. Therefore we can choose an increasing sequence n_k such that $n_k \in S_{\frac{1}{k}}$ for each $k \in \mathbb{N}$. Clearly $x_{n_k} \to x_\omega$ as $k \to \infty$.

Now assume that ω is selective. Consider the sets

$$C_k = \left\{\, n \in \mathbb{N} : \tfrac{1}{k} < |x_n - x_\omega| \leqslant \tfrac{1}{k-1} \,\right\},$$

where we assume $\frac{1}{0} = \infty$, and the set

$$C_\infty = \{\, n \in \mathbb{N} : x_n = x_\omega \,\}.$$

Note that $\omega(C_k) = 0$ for any $k \in \mathbb{N}$.

If $\omega(C_\infty) = 1$, we can take the subsequence consisting of the x_n, $n \in C_\infty$.

Otherwise, discarding all empty sets among C_k and C_∞ gives a partition of $\mathbb{N}$ into a countable collection of ω-small sets. Since ω is selective, we can choose a set $S \subset \mathbb{N}$ such that S meets each set of the partition at one point and $\omega(S) = 1$. Clearly the subsequence consisting of the x_n, $n \in S$ converges to x_ω in the usual sense. □

The following proposition is analogous to the statement that any sequence in a compact metric space has a convergent subsequence; it can be proved in the same way.

3.3. Proposition. *Let $\mathcal{X}$ be a compact metric space. Then any sequence of points x_n in $\mathcal{X}$ has a unique ω-limit x_ω.*

In particular, a bounded sequence of real numbers has a unique ω-limit.

The following lemma is an ultralimit analog of the Cauchy convergence test.

3.4. Lemma. *Let x_n be a sequence of points in a complete metric space $\mathcal{X}$. If for each subsequence y_n of x_n, the ω-limit*

$$y_\omega = \lim_{n\to\omega} y_n \in \mathcal{X}$$

is defined and does not depend on the choice of a subsequence, then the sequence x_n converges in the usual sense.

Proof. Assume the contrary. Then for some $\varepsilon > 0$, there is a subsequence y_n of x_n such that $|x_n - y_n| \geqslant \varepsilon$ for all n.

It follows that $|x_\omega - y_\omega| \geqslant \varepsilon$, a contradiction. □

C Ultralimits of spaces

Fix a selective ultrafilter ω on the set $\mathbb{N}$ of natural numbers.

Let $\mathcal{X}_n$ be a sequence of metric spaces. Consider all sequences $x_n \in \in \mathcal{X}_n$. On the set of all such sequences, define a pseudometric by the formula

➊ $$|(x_n) - (y_n)| = \lim_{n\to\omega} |x_n - y_n|.$$

Note that the ω-limit on the right-hand side is always defined and takes value in $[0, \infty]$.

Let $\mathcal{X}_\omega$ be the corresponding metric space; that is, the underlying set of $\mathcal{X}_\omega$ is formed by equivalence classes of sequences of points $x_n \in \mathcal{X}_n$ defined by

$$(x_n) \sim (y_n) \quad \Longleftrightarrow \quad \lim_{n\to\omega} |x_n - y_n| = 0,$$

and the distance is defined as in ➊.

The space $\mathcal{X}_\omega$ is called the ω-limit of $\mathcal{X}_n$. Typically $\mathcal{X}_\omega$ will denote the ω-limit of a sequence $\mathcal{X}_n$; we may also write

$$\mathcal{X}_n \to \mathcal{X}_\omega \quad \text{as} \quad n \to \omega \quad \text{or} \quad \mathcal{X}_\omega = \lim_{n\to\omega} \mathcal{X}_n.$$

Given a sequence $x_n \in \mathcal{X}_n$, we will denote by x_ω its equivalence class, which is a point in $\mathcal{X}_\omega$; equivalently, we will write

$$x_n \to x_\omega \quad \text{as} \quad n \to \omega \quad \text{or} \quad x_\omega = \lim_{n\to\omega} x_n.$$

3.5. Observation. *The* ω*-limit of any sequence of metric spaces is complete.*

Proof. Let $\mathcal{X}_n$ be a sequence of metric spaces and $\mathcal{X}_n \to \mathcal{X}_\omega$ as $n \to \omega$. Choose a Cauchy sequence x_n in $\mathcal{X}_\omega$. Passing to a subsequence, we can assume that $|x_k - x_m|_{\mathcal{X}_\omega} < \frac{1}{k}$ for any $k < m$.

Let us choose a double sequence $x_{n,m}$ in $\mathcal{X}_n$ such that for any fixed m we have $x_{n,m} \to x_m$ as $n \to \omega$. Note that for any $k < m$ the inequality $|x_{n,k} - x_{n,m}| < \frac{1}{k}$ holds for ω-almost all n. It follows that we can choose a nested sequence of sets

$$\mathbb{N} = S_1 \supset S_2 \supset \dots$$

such that

- $\omega(S_m) = 1$ for each m,
- $\bigcap_m S_m = \varnothing$, and
- $|x_{n,k} - x_{n,l}| < \frac{1}{k}$ for $k < l \leqslant m$ and $n \in S_m$.

Consider the sequence $y_n = x_{n,m(n)}$, where $m(n)$ is the largest value such that $n \in S_{m(n)}$. Denote by $y_\omega \in \mathcal{X}_\omega$ the ω-limit of y_n.

Observe that $|y_m - x_{n,m}| < \frac{1}{m}$ for ω-almost all n. It follows that $|x_m - y_\omega| \leqslant \frac{1}{m}$ for any m. Therefore, $x_n \to y_\omega$ as $n \to \infty$. That is, any Cauchy sequence in $\mathcal{X}_\omega$ converges. □

3.6. Observation. *The* ω*-limit of any sequence of length spaces is geodesic.*

Proof. If $\mathcal{X}_n$ is a sequence of length spaces, then for any sequence of pairs (x_n, y_n) in $\mathcal{X}_n$ there is a sequence of $\frac{1}{n}$-midpoints z_n.

Let $x_n \to x_\omega$, $y_n \to y_\omega$, and $z_n \to z_\omega$ as $n \to \omega$. Note that z_ω is a midpoint between x_ω and y_ω in $\mathcal{X}_\omega$.

By Observation 3.5, $\mathcal{X}_\omega$ is complete. Applying Lemma 2.12 we obtain the statement. □

Ultrapower. If all the metric spaces in a sequence are identical, $\mathcal{X}_n = \mathcal{X}$, the ω-limit $\lim_{n\to\omega} \mathcal{X}_n$ is denoted by $\mathcal{X}^\omega$ and called the ω-power of $\mathcal{X}$.

By Theorem 4.14, there is a distance-preserving map $\iota\colon \mathcal{X} \hookrightarrow \mathcal{X}^\omega$, where $\iota(y)$ is the equivalence class of the constant sequence $y_n = y$.

The image $\iota(\mathcal{X})$ might be a proper subset of $\mathcal{X}^\omega$. For example, $\mathbb{R}^\omega$ has pairs of points at distance ∞ from each other, while each metric component of $\mathbb{R}^\omega$ is isometric to $\mathbb{R}$.

According to Theorem 4.14, if $\mathcal{X}$ is compact then $\iota(\mathcal{X}) = \mathcal{X}^\omega$; in particular, $\mathcal{X}^\omega$ is isometric to $\mathcal{X}$. If $\mathcal{X}$ is proper, then $\iota(\mathcal{X})$ forms a metric component of $\mathcal{X}^\omega$.

The embedding ι allows us to treat $\mathcal{X}$ as a subset of its ultrapower $\mathcal{X}^\omega$.

3.7. Observation. *Let $\mathcal{X}$ be a complete metric space. Then $\mathcal{X}^\omega$ is a geodesic space if and only if $\mathcal{X}$ is a length space.*

Proof. Assume $\mathcal{X}^\omega$ is geodesic space. Then any pair of points $x, y \in \mathcal{X}$ has a midpoint $z_\omega \in \mathcal{X}^\omega$. Fix a sequence of points $z_n \in \mathcal{X}$ such that $z_n \to z_\omega$ as $n \to \omega$.

Note that $|x - z_n|_{\mathcal{X}} \to \frac{1}{2} \cdot |x - y|_{\mathcal{X}}$ and $|y - z_n|_{\mathcal{X}} \to \frac{1}{2} \cdot |x - y|_{\mathcal{X}}$ as $n \to \omega$. In particular, for any $\varepsilon > 0$, the point z_n is an ε-midpoint between x and y for ω-almost all n. It remains to apply Lemma 2.12.

The "if"-part follows from Observation 3.6. □

Note that the proof above together with Lemma 3.4 imply the following:

3.8. Corollary. *Assume $\mathcal{X}$ is a complete length space and $p, q \in \mathcal{X}$ cannot be joined by a geodesic in $\mathcal{X}$. Then there are at least two distinct geodesics between p and q in the ultrapower $\mathcal{X}^\omega$.*

D Ultralimits of sets

Let $\mathcal{X}_n$ be a sequence of metric spaces and $\mathcal{X}_n \to \mathcal{X}_\omega$ as $n \to \omega$.

For a sequence of sets $\Omega_n \subset \mathcal{X}_n$, consider the maximal set $\Omega_\omega \subset \mathcal{X}_\omega$ such that for any $x_\omega \in \Omega_\omega$ and any sequence $x_n \in \mathcal{X}_n$ such that $x_n \to x_\omega$ as $n \to \omega$, we have $x_n \in \Omega_n$ for ω-almost all n.

The set Ω_ω is called the open ω-limit of Ω_n; we could also write $\Omega_n \to \Omega_\omega$ as $n \to \omega$ or $\Omega_\omega = \lim_{n\to\omega} \Omega_n$.

Applying Observation 3.5 to the sequence of complements $\mathcal{X}_n \backslash \Omega_n$, we see that Ω_ω is open for any sequence Ω_n. The definition can be applied for arbitrary sequences of sets, but open ω-convergence will be applied here only for sequences of open sets.

E Ultralimits of functions

Recall that a family of submaps (see section 5A) between metric spaces $\{f_\alpha \colon \mathcal{X} \hookrightarrow \mathcal{Y}\}_{\alpha \in \mathcal{A}}$ is called equicontinuous if for any $\varepsilon > 0$ there is $\delta > 0$ such that for any $p, q \in \mathcal{X}$ with $|p - q| < \delta$ and any $\alpha \in \mathcal{A}$ we have $|f(p) - f(q)| < \varepsilon$.

Let $f_n \colon \mathcal{X}_n \hookrightarrow \mathbb{R}$ be a sequence of subfunctions.

Set $\Omega_n = \operatorname{Dom} f_n$. Consider the open ω-limit set $\Omega_\omega \subset \mathcal{X}_\omega$ of Ω_n.

Assume there is a subfunction $f_\omega \colon \mathcal{X}_\omega \hookrightarrow \mathbb{R}$ that satisfies the following conditions: (1) $\operatorname{Dom} f_\omega = \Omega_\omega$, (2) if $x_n \to x_\omega \in \Omega_\omega$ for a sequence of points $x_n \in \mathcal{X}_n$, then $f_n(x_n) \to f_\omega(x_\omega)$ as $n \to \omega$. In this case, the subfunction $f_\omega \colon \mathcal{X}_\omega \to \mathbb{R}$ is said to be the ω-limit of $f_n \colon \mathcal{X}_n \to \mathbb{R}$.

The following lemma gives a mild condition on a sequence of functions f_n guaranteeing the existence of its ω-limit.

3.9. Lemma. *Let $\mathcal{X}_n$ be a sequence of metric spaces and $f_n\colon \mathcal{X}_n \hookrightarrow \mathbb{R}$ be a sequence of subfunctions.*

Assume that for any positive integer $\Bbbk$, there is an open set $\Omega_n(\Bbbk) \subset \subset \mathrm{Dom}\, f_n$ such that the restrictions $f_n|_{\Omega_n(\Bbbk)}$ are uniformly bounded and equicontinuous and the open ω-limit of $\Omega_n(n)$ coincides with the open ω-limit of $\mathrm{Dom}\, f_n$. Then the ω-limit f_ω of f_n is defined; moreover f_ω is a continuous subfunction.

In particular, if the functions f_n are uniformly bounded and equicontinuous, then its ω-limit f_ω is defined, bounded and uniformly continuous.

The proof is straightforward.

3.10. Exercise. *Construct a sequence of compact length spaces $\mathcal{X}_n$ with a converging sequence of ℓ-Lipschitz concave (see Definition 5.17) functions $f_n\colon \mathcal{X}_n \to \mathbb{R}$ such that the ω-limit $\mathcal{X}_\omega$ of $\mathcal{X}_n$ is compact and the ω-limit $f_\omega\colon \mathcal{X}_\omega \to \mathbb{R}$ of f_n is not concave.*

If $f\colon \mathcal{X} \hookrightarrow \mathbb{R}$ is a subfunction, the ω-limit of the constant sequence $f_n = f$ is called the ω-power of f and is denoted by f^ω. So

$$f^\omega\colon \mathcal{X} \hookrightarrow \mathbb{R}, \quad f^\omega(x_\omega) = \lim_{n\to\omega} f(x_n).$$

Recall that we treat $\mathcal{X}$ as a subset of its ω-power $\mathcal{X}^\omega$. Note that $\mathrm{Dom}\, f = \mathcal{X} \cap \mathrm{Dom}\, f^\omega$. Moreover, if $\mathrm{B}(x,\varepsilon)_{\mathcal{X}} \subset \mathrm{Dom}\, f$ then $\mathrm{B}(x,\varepsilon)_{\mathcal{X}^\omega} \subset \subset \mathrm{Dom}\, f^\omega$.

Chapter 4

Space of spaces

In this chapter we discuss the Gromov–Hausdorff convergence of metric spaces.

To the best of our knowledge, Hausdorff convergence of subsets of a fixed metric space was first introduced by Felix Hausdorff [66], and a couple of years later an equivalent definition was given by Wilhelm Blaschke [22]. A further refinement of this definition was introduced by Zdeněk Frolík [52] and then rediscovered by Robert Wijsman [129]. However this refinement was a step in the direction of the so-called closed convergence introduced by Hausdorff in the original book. For that reason we call it Hausdorff convergence instead of Hausdorff–Blashcke–Frolík–Wijsman convergence.

Gromov–Hausdorff convergence was first introduced by David Edwards [48] and rediscovered later by Michael Gromov [57]. It was an essential tool in Gromov's proof that any group of polynomial growth has a nilpotent subgroup of finite index. Other versions of convergence of metric spaces were considered earlier, but each time the definition was limited to very specific types of problems.

The definition of Gromov–Hausdorff convergence of metric spaces uses the notion of Hausdorff convergence. Gromov–Hausdorff convergence means that a sequence of metric spaces admits a sequence of distance-preserving embeddings into a common ambient metric space so that their images converge in the Hausdorff sense. Our definition of Gromov–Hausdorff convergence and Gromov–Hausdorff distance differ somewhat from the standard definition.

A Convergence of subsets

Let $\mathcal{X}$ be a metric space and $A \subset \mathcal{X}$. Recall that the distance from A to a point x in $\mathcal{X}$ is given by

$$\operatorname{dist}_A x := \inf \{\, |a - x| : a \in A \,\}.$$

By this definition, we have $\operatorname{dist}_\varnothing x = \infty$ for any x.

4.1. Definition of Hausdorff convergence. *Given a sequence of closed sets A_n in a metric space $\mathcal{X}$, a closed set $A_\infty \subset \mathcal{X}$ is called the Hausdorff limit of A_n, briefly $A_n \xrightarrow[\text{H}]{} A_\infty$, if*

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

for any fixed $x \in \mathcal{X}$.

In this case the, sequence of closed sets A_n is said to be converging or converging in the sense of Hausdorff.

4.2. Selection theorem. *Let $\mathcal{X}$ be a proper metric space and A_n be a sequence of closed sets in $\mathcal{X}$. Then A_n has a converging subsequence in the sense of Hausdorff.*

Proof. Since $\mathcal{X}$ is proper, we can choose a countable dense set $\{x_1, x_2, \dots\}$ in $\mathcal{X}$.

If the sequence $a_n = \operatorname{dist}_{A_n} x_k$ is unbounded for some k, then we can pass to a subsequence of A_n such that $\operatorname{dist}_{A_n} x_k \to \infty$ as $n \to \infty$ for any k. The obtained sequence converges to the empty set.

Now suppose that $a_n = \operatorname{dist}_{A_n} x_k$ is bounded for each k. In this case, passing to a subsequence of A_n, 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 $f : \mathcal{X} \to \mathbb{R}$.

Set $A_\infty = f^{-1}(0)$. Since f is 1-Lipschitz, $\operatorname{dist}_{A_\infty} y \geqslant f(y)$ for any $y \in \mathcal{X}$. It remains to show that $\operatorname{dist}_{A_\infty} y \leqslant f(y)$ for any y.

Assume the contrary, that is, $f(z) < R < \operatorname{dist}_{A_\infty} z$ 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 $n \to \infty$.

Clearly $f(z_\infty) = 0$, that is, $z_\infty \in A_\infty$. On the other hand,

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

a contradiction. ☐

B Convergence of spaces

4.3. Definition. *Let* $\{\mathcal{X}_\alpha : \alpha \in \mathcal{A}\}$ *be a set of metric spaces. A metric space* $\boldsymbol{X}$ *is called a common space of* $\{\mathcal{X}_\alpha : \alpha \in \mathcal{A}\}$ *if its underlying set is formed by the disjoint union*

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

and each inclusion $\iota_\alpha\colon \mathcal{X}_\alpha \hookrightarrow \boldsymbol{X}$ *is distance-preserving.*

4.4. Definition. *Let* $\boldsymbol{X}$ *be a common space for proper metric spaces* $\mathcal{X}_1, \mathcal{X}_2, \dots,$ *and* $\mathcal{X}_\infty$*. Assume that* $\mathcal{X}_n$ *forms an open set in* $\boldsymbol{X}$ *for each* $n < \infty$ *and* $\mathcal{X}_n \xrightarrow[\mathrm{H}]{} \mathcal{X}_\infty$ *in* $\boldsymbol{X}$ *as* $n \to \infty$*.*

Then the topology τ *of* $\boldsymbol{X}$ *is called a Gromov–Hausdorff convergence and we write* $\mathcal{X}_n \xrightarrow[\mathrm{GH}]{} \mathcal{X}_\infty$ *or* $\mathcal{X}_n \xrightarrow[\mathrm{GH}]{\tau} \mathcal{X}_\infty$*; the latter notation is used if we need to consider the specific Gromov–Hausdorff convergence* τ*. The space* $\mathcal{X}_\infty$ *is called the limit space of the sequence* $\mathcal{X}_n$ *along* τ*.*

When we write $\mathcal{X}_n \xrightarrow[\mathrm{GH}]{} \mathcal{X}_\infty$ we mean that we made a choice of a Gromov–Hausdorff convergence.

Note that for a fixed sequence $\mathcal{X}_n$ of metric spaces, one may construct different Gromov–Hausdorff convergences, say $\mathcal{X}_n \xrightarrow[\mathrm{GH}]{\tau} \mathcal{X}_\infty$ and $\mathcal{X}_n \xrightarrow[\mathrm{GH}]{\tau'} \xrightarrow[\mathrm{GH}]{\tau'} \mathcal{X}'_\infty$, and their limit spaces $\mathcal{X}_\infty$ and $\mathcal{X}'_\infty$ need not be isometric to each other. For example, for the constant sequence $\mathcal{X}_n \overset{iso}{=\!=} \mathbb{R}_{\geqslant 0}$, one may take $\mathcal{X}_\infty \overset{iso}{=\!=} \mathbb{R}_{\geqslant 0}$. In this case, a point in the disjoint space $\boldsymbol{X}$ can be regarded as a pair $(x, n) \in \mathbb{R}_{\geqslant 0} \times (\mathbb{Z}_{>} \cup \{\infty\})$ and the metric on $\boldsymbol{X}$ can be defined by

$$|(x,n) - (y,m)|_{\boldsymbol{X}} := |\tfrac{1}{n} - \tfrac{1}{m}| + |x - y|,$$

where we assume that $0 = \frac{1}{\infty}$. On the other hand, one can take $\mathcal{X}'_\infty \overset{iso}{=\!=} \mathbb{R}$, and consider the metric

$$\begin{aligned} |(x,n) - (y,m)|_{\boldsymbol{X}'} &= |\tfrac{1}{n} - \tfrac{1}{m}| + |(x-n) - (y-m)|, \\ |(x,n) - (y,\infty)|_{\boldsymbol{X}'} &= \tfrac{1}{n} + |(x-n) - y|, \\ |(x,\infty) - (y,\infty)|_{\boldsymbol{X}'} &= |x - y|. \end{aligned}$$

where $n, m < \infty$

4.5. Induced convergences. *Suppose* $\mathcal{X}_n \xrightarrow[\mathrm{GH}]{\tau} \mathcal{X}_\infty$ *as in Definition 4.4, and* $\iota_n\colon \mathcal{X}_n \hookrightarrow \boldsymbol{X}$*,* $\iota_\infty\colon \mathcal{X}_\infty \hookrightarrow \boldsymbol{X}$ *are the corresponding inclusions.*

a) A sequence of points $x_n \in \mathcal{X}_n$ *converges to* $x_\infty \in \mathcal{X}_\infty$ *(briefly,* $x_n \to x_\infty$ *or* $x_n \xrightarrow{\tau} x_\infty$*) if* $|x_n - x_\infty|_{\boldsymbol{X}} \to 0$*.*

b) *A sequence of closed sets* $\mathfrak{C}_n \subset \mathcal{X}_n$ *converges to a closed set* $\mathfrak{C}_\infty \subset \mathcal{X}_\infty$ *(briefly,* $\mathfrak{C}_n \to \mathfrak{C}_\infty$ *or* $\mathfrak{C}_n \xrightarrow{\tau} \mathfrak{C}_\infty$*) if* $\mathfrak{C}_n \xrightarrow[\mathrm{H}]{} \mathfrak{C}_\infty$ *as subsets of* $\boldsymbol{X}$.

c) *A sequence of open sets* $\Omega_n \subset \mathcal{X}_n$ *converges to an open set* $\Omega_\infty \subset \mathcal{X}_\infty$ *(briefly,* $\Omega_n \to \Omega_\infty$ *or* $\Omega_n \xrightarrow{\tau} \Omega_\infty$*) if the complements* $\mathcal{X}_n \backslash \Omega_n$ *converge to the complement* $\mathcal{X}_\infty \backslash \Omega_\infty$ *as closed sets.*

d) *Let* $\mathcal{X}_n \xrightarrow[\mathrm{GH}]{\tau} \mathcal{X}_\infty$ *and* $\mathcal{Y}_n \xrightarrow[\mathrm{GH}]{\vartheta} \mathcal{Y}_\infty$. *A sequence of submaps (see section 5A)* $\Phi_n \colon \mathcal{X}_n \hookrightarrow \mathcal{Y}_n$ *converges to a submap* $\Phi_\infty \colon \mathcal{X}_\infty \hookrightarrow \mathcal{Y}_\infty$ *if the following conditions holds*

- $\mathrm{Dom}\,\Phi_n \to \mathrm{Dom}\,\Phi_\infty$ *as a sequence of open sets.*
- *for any* $x_\infty \in \mathrm{Dom}\,\Phi_\infty$ *and any sequence* $x_n \in \mathcal{X}_n$ *such that* $x_n \to x_\infty$,

$$\mathcal{Y}_n \ni \Phi_n(x_n) \xrightarrow{\vartheta} \Phi_\infty(x_\infty) \in \mathcal{Y}_\infty$$

as $n \to \infty$.

e) *Given a sequence of measures* μ_n *on* $\mathcal{X}_n$, *we say that* μ_n *weakly converges to a measure* μ_∞ *on* $\mathcal{X}_\infty$ *(briefly,* $\mu_n \to \mu_\infty$ *or* $\mu_n \xrightarrow{\tau} \mu_\infty$*) if the pushforward measures of* μ_n *weakly converge to the pushforward measure of* μ_∞.
In other words, if for any continuous function $\varphi \colon \boldsymbol{X} \to \mathbb{R}$ *with a compact support, we have*

$$\int\limits_{\mathcal{X}_n} \varphi \circ \iota_n(x)\cdot \boldsymbol{d}_x\mu_n \to \int\limits_{\mathcal{X}_\infty} \varphi \circ \iota_\infty(x)\cdot \boldsymbol{d}_x\mu_\infty$$

as $n \to \infty$.

Liftings. Given a Gromov–Hausdorff convergence $\mathcal{X}_n \xrightarrow[\mathrm{GH}]{} \mathcal{X}_\infty$ and a point $p_\infty \in \mathcal{X}_\infty$, any sequence of points $p_n \in \mathcal{X}_n$ such that $p_n \xrightarrow[\mathrm{GH}]{} p_\infty$ will be called a lifting of p_∞. The point $p_n \in \mathcal{X}_n$ will be called a lifting of p_∞ in $\mathcal{X}_n$. We will also say that $\mathrm{dist}_{p_n} : \mathcal{X}_n \to \mathbb{R}$ is a lifting of the distance function $\mathrm{dist}_{p_\infty} : \mathcal{X}_\infty \to \mathbb{R}$. Clearly $\mathrm{dist}_{p_n} \xrightarrow[\mathrm{GH}]{} \mathrm{dist}_{p_\infty}$.

Note that liftings are not uniquely defined. We will be interested in the properties of liftings for sufficiently large n.

Similarly, we may refer to liftings of a point array $\boldsymbol{p}_\infty = (p^1_\infty, p^2_\infty, \dots$ $\dots, p^{\bar{k}}_\infty)$ and of the corresponding distance map $\mathrm{dist}_{\boldsymbol{p}_\infty} : \mathcal{X}_\infty \to \mathbb{R}^{\bar{k}}$,

$$\mathrm{dist}_{\boldsymbol{p}_\infty} : x \mapsto (|p^1_\infty - x|, |p^2_\infty - x|, \dots, |p^k_\infty - x|).$$

C Gromov's selection theorem

4.6. Gromov's 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 $R > 0$, $\varepsilon > 0$, there is $N = N(R, \varepsilon) \in \mathbb{Z}_{>0}$ such that for each n the ball $\overline{\mathrm{B}}[x_n, R] \subset \mathcal{X}_n$ admits a finite ε-net with at most N points. Then there is a subsequence of $\mathcal{X}_n$ that admits a Gromov–Hausdorff convergence such that the sequence of marked points $x_n \in \mathcal{X}_n$ converges.*

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

$$|z_{i,n} - x_n|_{\mathcal{X}_n} \leqslant k + 1 \quad \text{if} \quad i \leqslant M_k$$

and the points $z_{1,n}, \dots, z_{M_k,n}$ form an $\frac{1}{k}$-net in $\overline{\mathrm{B}}[x_n, k]_{\mathcal{X}_n}$.

Passing to a subsequence, we may assume that the sequence

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

converges for any i and j.

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

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

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

satisfying definitions 4.3 and 4.4.

Such a metric can be defined as follows. Fix a sequence $\varepsilon_k \to 0+$ and let N_k be the minimal integer such that

$$\big||w_i - w_j|_{\mathcal{W}} - |z_{i,n} - z_{j,n}|_{\mathcal{X}_n}\big| < \varepsilon_k$$

if $i, j \leqslant N_k$ and $n \geqslant N_k$. Let us equip $\boldsymbol{X}$ with the maximal metric such that all the inclusions $\iota_n \colon \mathcal{X}_n \to \boldsymbol{X}$ and $\iota_\infty \colon \mathcal{X}_\infty \to \boldsymbol{X}$ are isometric and

$$|z_{i,n} - w_i| \leqslant \varepsilon_k$$

for $i \leqslant N_k, n \geqslant N_k$. It is easy to verify that such a metric on $\boldsymbol{X}$ satisfies 4.3 and 4.4.

□

D Convergence of compact spaces.

4.7. Definition. *Let $\mathcal{X}$ and $\mathcal{Y}$ be metric spaces. A map $f\colon \mathcal{X} \to \mathcal{Y}$ is called an* ε-*isometry if the following two conditions hold:*

a) $\operatorname{Im} f$ *is an* ε*-net in* $\mathcal{Y}$.

b) $||f(x) - f(x')|_{\mathcal{Y}} - |x - x'|_{\mathcal{X}}| \leqslant \varepsilon$ *for any* $x, x' \in \mathcal{X}$.

4.8. Lemma. *Let $\mathcal{X}_1, \mathcal{X}_2, \dots$, and $\mathcal{X}_\infty$ be metric spaces and let $\varepsilon_n \to 0+$ as $n \to \infty$. Suppose that either*

a) for each n there is an ε_n-isometry $f_n\colon \mathcal{X}_n \to \mathcal{X}_\infty$, or

b) for each n there is an ε_n-isometry $h_n\colon \mathcal{X}_\infty \to \mathcal{X}_n$.

Then there is a Gromov–Hausdorff convergence $\mathcal{X}_n \xrightarrow[\text{GH}]{} \mathcal{X}_\infty$.

Proof. To prove part (4.8*a*) let us construct a common space $\boldsymbol{X}$ for the spaces $\mathcal{X}_1, \mathcal{X}_2, \dots$, and $\mathcal{X}_\infty$ by taking the metric ρ on the disjoint union $\mathcal{X}_\infty \sqcup \mathcal{X}_1 \sqcup \mathcal{X}_2 \sqcup \dots$ that is defined the following way:

$$\begin{aligned}
\rho(x_n, y_n) &= |x_n - y_n|_{\mathcal{X}_n}, \quad \rho(x_\infty, y_\infty) = |x_\infty - y_\infty|_{\mathcal{X}_\infty},\\
\rho(x_n, x_\infty) &= \inf\left\{\, |x_n - y_n|_{\mathcal{X}_n} + \varepsilon_n + |x_\infty - f(y_n)|_{\mathcal{X}_\infty} : y_n \in \mathcal{X}_n \,\right\},\\
\rho(x_n, x_m) &= \inf\left\{\, \rho(x_n, y_\infty) + \rho(x_m, y_\infty) : y_\infty \in \mathcal{X}_\infty \,\right\},
\end{aligned}$$

where we assume that $x_m \in \mathcal{X}_m$, $x_n \in \mathcal{X}_n$, and $x_\infty \in \mathcal{X}_\infty$.

It remains to observe that ρ is indeed a metric and $\mathcal{X}_n \xrightarrow[\text{H}]{} \mathcal{X}_\infty$ in $\boldsymbol{X}$.

The proof of the second part is analogous; one only needs to change one line in the definition of ρ to the following:

$$\rho(x_n, x_\infty) = \inf\left\{\, |x_n - h(y_\infty)|_{\mathcal{X}_n} + \varepsilon_n + |x_\infty - y_\infty|_{\mathcal{X}_\infty} : y_\infty \in \mathcal{X}_\infty \,\right\}.$$

□

4.9. Definition. *Given two compact spaces $\mathcal{X}$ and $\mathcal{Y}$, we will write*

- ◇ *$\mathcal{X} \leqslant \mathcal{Y}$ if there is a noncontracting map $\Phi\colon \mathcal{X} \to \mathcal{Y}$.*
- ◇ *$\mathcal{X} \leqslant \mathcal{Y} + \varepsilon$ if there is a map $\Phi\colon \mathcal{X} \to \mathcal{Y}$ such that for any $x, x' \in \mathcal{X}$ we have*

$$|x - x'| \leqslant |\Phi(x) - \Phi(x')| + \varepsilon.$$

4.10. Lemma. *Let $\mathcal{X}$ and $\mathcal{Y}$ be two metric spaces and $\mathcal{X}$ be compact. Then*

$$\mathcal{X} \geqslant \mathcal{Y} \geqslant \mathcal{X} \quad \Longleftrightarrow \quad \mathcal{X} \overset{iso}{=\!=} \mathcal{Y}.$$

The following proof was suggested by Travis Morrison.

Proof. Let $f\colon \mathcal{X} \to \mathcal{Y}$ and $g\colon \mathcal{Y} \to \mathcal{X}$ be noncontracting mappings. It is sufficient to prove that $h = g \circ f\colon \mathcal{X} \to \mathcal{X}$ is an isometry.

Given any pair of points $x, y \in \mathcal{X}$, set $x_n = h^{\circ n}(x)$ and $y_n = h^{\circ n}(y)$, where $\circ n$ denotes iteration n times. Since $\mathcal{X}$ is compact, one can choose an increasing sequence of integers n_k such that both sequences x_{n_i} and y_{n_i} converge. In particular, both of these sequences are Cauchy, that is,

$$|x_{n_i} - x_{n_j}|, |y_{n_i} - y_{n_j}| \to 0$$

as $\min\{i, j\} \to \infty$. Since h is noncontracting, we have

$$|x - x_{|n_i - n_j|}| \leqslant |x_{n_i} - x_{n_j}|.$$

It follows that there is a sequence $m_i \to \infty$ such that

❶ $$x_{m_i} \to x \quad \text{and} \quad y_{m_i} \to y \quad \text{as} \quad k \to \infty.$$

Let $\ell_n = |x_n - y_n|$. Since h is noncontracting, the sequence ℓ_n is nondecreasing. On the other hand, from ❶ it follows that $\ell_{m_i} \to |x - y| = \ell_0$ as $m_i \to \infty$, that is, the sequence ℓ_n is constant. In particular $\ell_0 = \ell_1$ for any x and y in $\mathcal{X}$, so h is a distance-preserving map.

Thus $h(X)$ is isometric to $\mathcal{X}$. From ❶, $h(\mathcal{X})$ is everywhere dense. Since X is compact, $h(\mathcal{X}) = \mathcal{X}$. □

The Gromov–Hausdorff distance between isometry classes of compact metric spaces $\mathcal{X}$ and $\mathcal{Y}$, is defined by

$$\mathrm{D_{GH}}(\mathcal{X}, \mathcal{Y}) := \inf \{ \varepsilon > 0 : \mathcal{X} \leqslant \mathcal{Y} + \varepsilon \text{ and } \mathcal{Y} \leqslant \mathcal{X} + \varepsilon \}.$$

The Gromov–Hausdorff distance turns the set of all isometry classes of compact metric spaces into a metric space. The following theorem shows that convergence in this space coincides with the Gromov–Hausdorff convergence defined above.

4.11. Theorem. *Let $\mathcal{X}_1, \mathcal{X}_2, \ldots,$ and $\mathcal{X}_\infty$ be compact metric spaces. Then there is a convergence $\mathcal{X}_n \xrightarrow[\mathrm{GH}]{} \mathcal{X}_\infty$ if and only if $\mathrm{D_{GH}}(\mathcal{X}_n, \mathcal{X}_\infty) \to 0$ as $n \to \infty$.*

Proof; if part. Suppose $a_n \colon \mathcal{X}_\infty \to \mathcal{X}_n$ and $b_n \colon \mathcal{X}_n \to \mathcal{X}_\infty$ are sequences of maps such that

$$|a_n(x) - a_n(y)|_{\mathcal{X}_\infty} \geqslant |x - y|_{\mathcal{X}_n} - \delta_n,$$

$$|b_n(v) - b_n(w)|_{\mathcal{X}_n} \geqslant |v - w|_{\mathcal{X}_\infty} - \delta_n$$

for any $x, y \in \mathcal{X}_n$ and $v, w \in \mathcal{X}_\infty$, and some sequence $\delta_n \to 0+$. Let us show that $\mathcal{X}_n \xrightarrow{\tau} \mathcal{X}_\infty$.

Fix $\varepsilon > 0$ and choose a maximal ε-packing $\{x^1, x^2, \ldots, x^k\}$ in $\mathcal{X}_\infty$ such that $\sum_{i<j} |x^i - x^j|$ is maximal. Note that

$$|a_n \circ b_n(x^i) - a_n \circ b_n(x^j)| \geqslant |x^i - x^j| - 2{\cdot}\delta_n.$$

Since $\sum_{i<j}|x^i - x^j|$ is maximal,

$$|a_n \circ b_n(x^i) - a_n \circ b_n(x^j)| \to |x^i - x^j|$$

for all i and j as $n \to \infty$. For all large n, we have $2\cdot\delta_n < |x^i - x^j| - \varepsilon$, and so

$$|b_n(x^i) - b_n(x^j)|_{\mathcal{X}_n} > \varepsilon \quad \text{and} \quad |a_n \circ b_n(x^i) - a_n \circ b_n(x^j)|_{\mathcal{X}_n} > \varepsilon$$

for all $i \neq j$. Therefore for each large n, the set $\{a_n \circ b_n(x^i)\}$ forms a maximal ε-packing and hence an ε-net in $\mathcal{X}_\infty$.

Since $\{a_n \circ b_n(x^i)\}$ is an ε-net in $\mathcal{X}_\infty$, then for any $y_n \in \mathcal{X}_n$ there is x^i such that $|a_n \circ b_n(x^i) - a_n(y_n)| < \varepsilon$. Thus $|b_n(x^i) - y_n| < \varepsilon + \delta_n$, that is, $\{b_n(x^i)\}$ is a $(\varepsilon + \delta_n)$-net in $\mathcal{X}_n$.

Given $y \in \mathcal{X}_n$, choose x^i so that $|b_n(x^i) - y_n| < \varepsilon + \delta_n$ and define $h_n(y) = a_n \circ b_n(x^i)$. Observe that h_n is a $3\cdot\varepsilon$-isometry for all large n. Since $\varepsilon > 0$ is arbitrary, there is a sequence of ε_n-isometries $\mathcal{X}_n \to \mathcal{X}_\infty$ such that $\varepsilon_n \to 0+$ as $n \to \infty$. It remains to apply 4.8.

Only-if part. Assume $\mathcal{X}_n \xrightarrow{\tau} \mathcal{X}_\infty$. Fix $\varepsilon > 0$, and choose a maximal ε-packing $\{x^1, x^2, \dots, x^k\}$ in $\mathcal{X}_\infty$. For each x^i, choose a sequence $x^i_n \in \mathcal{X}_n$ such that $a_n(x^i_n) \to x^i$. Define a map $a_n\colon \mathcal{X}_n \to \mathcal{X}_\infty$ such that $a_n(x^i_n) = x_n$. Note that for all large n, we have $|x^i_n - x^j_n| > \varepsilon$. For each point $z \in \mathcal{X}_\infty$, choose x^i so that $|z - x^i| < \varepsilon$. Define a map $b_n\colon \mathcal{X}_\infty \to \mathcal{X}_n$ such that $b_n(z) = x^i_n$. Observe that

$$|b_n(y) - b_n(z)|_{\mathcal{X}_n} + 3\cdot\varepsilon > |y - z|_{\mathcal{X}_\infty}$$

for all large n.

In the same way we can construct a map $a_n\colon \mathcal{X}_n \to \mathcal{X}_\infty$ such that

$$|a_n(y) - a_n(z)|_{\mathcal{X}_\infty} + 3\cdot\varepsilon > |y - z|_{\mathcal{X}_n}.$$

Hence $\mathrm{D_{GH}}(\mathcal{X}_n, \mathcal{X}_\infty) \to 0$ as $n \to \infty$. ☐

The following theorem states that the isometry class of a Gromov–Hausdorff limit is uniquely defined if it is compact.

4.12. Theorem. *Let $\mathcal{X}_1, \mathcal{X}_2, \dots$, and $\mathcal{X}_\infty$ and $\bar{\mathcal{X}}_\infty$ be metric spaces such that $\mathcal{X}_n \xrightarrow{\tau} \mathcal{X}_\infty$, $\mathcal{X}_n \xrightarrow{\bar{\tau}} \bar{\mathcal{X}}_\infty$.*

Assume that $\bar{\mathcal{X}}_\infty$ is compact. Then $\mathcal{X}_\infty \overset{iso}{=\!=} \bar{\mathcal{X}}_\infty$.

Proof. For each point $x_\infty \in \mathcal{X}_\infty$, choose liftings $x_n \in \mathcal{X}_n$.

Choose a nonprincipal ultrafilter ω on $\mathbb{N}$. Define $\bar{x}_\infty \in \bar{\mathcal{X}}_\infty$ as the ω-limit of x_n with respect to $\bar{\tau}$. We claim that the map $x_\infty \to \bar{x}_\infty$ is an isometry.

Indeed, by the definition of Gromov–Hausdorff convergence,

$$|\bar{x}_\infty - \bar{y}_\infty|_{\bar{\mathcal{X}}_\infty} = \lim_{n\to\omega} |x_n - y_n|_{\mathcal{X}_n} = |x_\infty - y_\infty|_{\mathcal{X}_\infty}.$$

Thus the map $x_\infty \to \bar{x}_\infty$ gives a distance-preserving map $\Phi\colon \mathcal{X}_\infty \hookrightarrow \bar{\mathcal{X}}_\infty$. In particular, $\mathcal{X}_\infty$ is compact. Switching $\mathcal{X}_\infty$ and $\bar{\mathcal{X}}_\infty$ and applying the same argument, we get an isometric embedding $\bar{\mathcal{X}}_\infty \hookrightarrow \mathcal{X}_\infty$. Now the result follows from Lemma 4.10. ☐

4.13. Exercise. *Let $\mathcal{X}_n$ be a sequence of metric spaces that admits two Gromov–Hausdorff convergences τ and τ'. Assume $\mathcal{X}_n \xrightarrow[\text{GH}]{\tau} \mathcal{X}_\infty$ and $\mathcal{X}_n \xrightarrow[\text{GH}]{\tau'} \mathcal{X}'_\infty$. Show that if $\mathcal{X}_\infty$ is proper and there is a sequence of points $x_n \in \mathcal{X}_n$ that converges in both τ and τ', then $\mathcal{X}_\infty \overset{iso}{=\!=} \mathcal{X}'_\infty$.*

E Gromov–Hausdorff convergence and ultralimits

4.14. Theorem. *Assume $\mathcal{X}_n$ is a sequence of complete metric spaces. Let $\mathcal{X}_n \to \mathcal{X}_\omega$ as $n \to \omega$, and let $\mathcal{Y}_n$ be a sequence of subsets of the $\mathcal{X}_n$ such that $\mathcal{Y}_n \xrightarrow[\text{GH}]{} \mathcal{Y}_\infty$. Then there is a distance-preserving map $\iota : \mathcal{Y}_\infty \to \mathcal{X}_\omega$.*

Moreover:

a) If $\mathcal{X}_n \xrightarrow[\text{GH}]{} \mathcal{X}_\infty$ and $\mathcal{X}_\infty$ is compact, then $\mathcal{X}_\infty$ is isometric to $\mathcal{X}_\omega$.

b) If $\mathcal{X}_n \xrightarrow[\text{GH}]{} \mathcal{X}_\infty$ and $\mathcal{X}_\infty$ is proper, then $\mathcal{X}_\infty$ is isometric to a metric component of $\mathcal{X}_\omega$.

Proof. For each point $y_\infty \in \mathcal{Y}_\infty$ choose a lifting $y_n \in \mathcal{Y}_n$. Pass to the ω-limit $y_\omega \in \mathcal{X}_\omega$ of y_n. Clearly for any $y_\infty, z_\infty \in \mathcal{Y}_\infty$, we have

$$|y_\infty - z_\infty|_{\mathcal{Y}_\infty} = |y_\omega - z_\omega|_{\mathcal{X}_\omega};$$

that is, the map $y_\infty \mapsto y_\omega$ gives a distance-preserving map $\iota : \mathcal{Y}_\infty \to \mathcal{X}_\omega$.

(a)+(b). Fix $x_\omega \in \mathcal{X}_\omega$. Choose a sequence x_n of elements of $\mathcal{X}_n$, such that $x_n \to x_\omega$ as $n \to \omega$.

Denote by $\boldsymbol{X} = \mathcal{X}_\infty \sqcup \mathcal{X}_1 \sqcup \mathcal{X}_2 \sqcup \ldots$ the common space for the convergence $\mathcal{X}_n \xrightarrow[\text{GH}]{} \mathcal{X}_\infty$, as in the definition of Gromov–Hausdorff convergence. Consider the sequence x_n as a sequence of points in $\boldsymbol{X}$.

If the ω-limit x_∞ of x_n in $\boldsymbol{X}$ exists, it must lie in $\mathcal{X}_\infty$.

The point x_∞, if defined, does not depend on the choice of x_n. Indeed, if $y_n \in \mathcal{X}_n$ is another sequence such that $y_n \to x_\omega$ as $n \to \omega$, then

$$|y_\infty - x_\infty| = \lim_{n\to\omega} |y_n - x_n| = 0;$$

that is, $x_\infty = y_\infty$.

In this way we obtain a map $\nu\colon x_\omega \to x_\infty$, defined on $\operatorname{Dom}\nu \subset \mathcal{X}_\omega$. By construction of ι, we have $\iota\circ\nu(x_\omega) = x_\omega$ for any $x_\omega \in \operatorname{Dom}\nu$.

Finally note that if $\mathcal{X}_\infty$ is compact, then ν is defined on all of $\mathcal{X}_\omega$; this proves (a).

If $\mathcal{X}_\infty$ is proper, choose any point $z_\infty \in \mathcal{X}_\infty$ and set $z_\omega = \iota(z_\infty)$. For any point $x_\omega \in \mathcal{X}_\omega$ at finite distance from z_ω, for the sequence x_n as above we have that $|z_n - x_n|$ is bounded for ω-almost all n. Since $\mathcal{X}_\infty$ is proper, $\nu(x_\omega)$ is defined; in other words, ν is defined on the metric component of z_ω. Hence (b) follows. $\square$

Chapter 5

Maps and functions

Here we introduce some classes of maps between metric spaces and develop a language to describe various notions of convexity/concavity of real-valued functions on general metric space.

A Submaps

We will often need maps and functions defined on subsets of a metric space. We call them submaps and subfunctions. Thus, given metric spaces $\mathcal{X}$ and $\mathcal{Y}$, a submap $\Phi\colon \mathcal{X} \multimapdotinv \mathcal{Y}$ is a map defined on a subset $\operatorname{Dom}\Phi \subset \mathcal{X}$.

A submap is said to be continuous if the inverse image of any open set is open. Note that for a continuous submap Φ, the domain $\operatorname{Dom}\Phi$ is automatically open. Indeed, if submap $\Phi\colon \mathcal{X} \multimapdotinv \mathcal{Y}$ is continuous then $\operatorname{Dom}\Phi = \Phi^{-1}(\mathcal{Y})$ is open as the preimage of an open set. The same is true for upper and lower semicontinuous functions $f\colon \mathcal{X} \multimapdotinv \mathbb{R}$ since they are continuous functions for a special topology on $\mathbb{R}$.

(Continuous partially defined maps could be defined via closed sets; namely, one could require that inverse images of closed sets are closed. While this condition is equivalent to continuity for functions defined on the whole space, it is different for partially defined functions. In particular, with this definition the domain of a continuous submap would have to be closed.)

B Lipschitz conditions

5.1. Lipschitz maps. *Suppose $\mathcal{X}$ and $\mathcal{Y}$ are metric spaces, $\Phi\colon \mathcal{X} \multimapdotinv \mathcal{Y}$ is a continuous submap, and $\ell \in \mathbb{R}$.*

a) The submap Φ *is called* ℓ*-Lipschitz if*

$$|\Phi(x) - \Phi(y)|_{\mathcal{Y}} \leqslant \ell\cdot|x - y|_{\mathcal{X}}$$

for any two points $x, y \in \operatorname{Dom}\Phi$*.*

- 1*-Lipschitz maps will be also called short.*

b) We say that Φ *is Lipschitz if it is* ℓ*-Lipschitz for some constant* ℓ*. The minimal such constant is denoted by* $\operatorname{lip}\Phi$*.*

c) We say that Φ *is locally Lipschitz if any point* $x \in \operatorname{Dom}\Phi$ *admits a neighborhood* $\Omega \subset \operatorname{Dom}\Phi$ *such that the restriction* $\Phi|_\Omega$ *is Lipschitz.*

d) Given $p \in \operatorname{Dom}\Phi$*, we denote by* $\operatorname{lip}_p\Phi$ *the infimum of the real values* ℓ *such that* p *admits a neighborhood* $\Omega \subset \operatorname{Dom}\Phi$ *such that the restriction* $\Phi|_\Omega$ *is* ℓ*-Lipschitz.*

Note that $\Phi\colon \mathcal{X} \to \mathcal{Y}$ is ℓ-Lipschitz if and only if

$$\Phi(\mathrm{B}(x, R)_{\mathcal{X}}) \subset \mathrm{B}(\Phi(x), \ell\cdot R)_{\mathcal{Y}}$$

for any $R \geqslant 0$ and $x \in \mathcal{X}$. The following definition gives a dual version.

5.2. Definitions. *Let* $\mathcal{X}$ *and* $\mathcal{Y}$ *be metric spaces,* $\Phi\colon \mathcal{X} \to \mathcal{Y}$ *be a map, and* $\ell \in \mathbb{R}$*.*

a) The map Φ *is called* ℓ*-co-Lipshitz if*

$$\Phi(\mathrm{B}(x, \ell\cdot R)_{\mathcal{X}}) \supset \mathrm{B}(\Phi(x), R)_{\mathcal{Y}}$$

for any $x \in \mathcal{X}$ *and* $R > 0$*.*

b) The map Φ *is called co-Lipschitz if it is* ℓ*-co-Lipschitz for some constant* ℓ*. The minimal such constant is denoted by* $\operatorname{colip}\Phi$*.*

From the definition of co-Lipschitz maps we get the following:

5.3. Proposition. *Any co-Lipschitz map is open and surjective.*

In other words, ℓ-co-Lipschitz maps can be considered as a quantitative version of open maps. For that reason they are also called ℓ-open [34]. Also, be aware that some authors refer to our ℓ-co-Lipschitz maps as $\frac{1}{\ell}$-co-Lipschitz.

5.4. Proposition. *Let* $\mathcal{X}$ *and* $\mathcal{Y}$ *be metric spaces such that* $\mathcal{X}$ *is complete, and let* $\Phi\colon \mathcal{X} \to \mathcal{Y}$ *be a continuous co-Lipschitz map. Then* $\mathcal{Y}$ *is complete.*

Proof. Choose a Cauchy sequence y_n in $\mathcal{Y}$. Passing to a subsequence if necessary, we may assume that $|y_n - y_{n+1}|_{\mathcal{Y}} < \frac{1}{2^n}$ for each n.

Denote by ℓ a co-Lipschitz constant for Φ. Note that there is a sequence x_n in $\mathcal{X}$ such that

❶ $$\Phi(x_n) = y_n \quad \text{and} \quad |x_n - x_{n+1}|_{\mathcal{X}} < \tfrac{\ell}{2^n}$$

for each n. Indeed, such a sequence can be constructed recursively. Assuming that the points $x_1, \dots, x_{n-1}$ are already constructed, the existence of a sequence x_n satisfying ❶ follows since Φ is ℓ-co-Lipschitz.

Notice that the sequence x_n is Cauchy. Since $\mathcal{X}$ is complete, x_n converges in $\mathcal{X}$; denote its limit by x_∞ and set $y_\infty = \Phi(x_\infty)$. Since Φ is continuous, $y_n \to y_\infty$ as $n \to \infty$. Hence the result. □

5.5. Lemma. *Let $\mathcal{X}$ be a metric space and $f\colon \mathcal{X} \to \mathbb{R}$ be a continuous function. Then for any $\varepsilon > 0$ there is a locally Lipschitz function $f_\varepsilon : \mathcal{X} \to \mathbb{R}$ such that $|f(x) - f_\varepsilon(x)| < \varepsilon$ for any $x \in \mathcal{X}$.*

Proof. Assume that $f \geqslant 1$. Construct a continuous positive function $\rho\colon \mathcal{X} \to \mathbb{R}_{>0}$ such that

$$|x - y| < \rho(x) \quad \Rightarrow \quad |f(x) - f(y)| < \varepsilon.$$

Consider the function

$$f_\varepsilon(x) = \sup\left\{\, f(x)\cdot \left(1 - \tfrac{|x-y|}{\rho(x)}\right) \;:\; x \in \mathcal{X} \,\right\}.$$

It is straightforward to check that each f_ε is locally Lipschitz and $0 \leqslant f_\varepsilon - f < \varepsilon$.

Since any continuous function can be presented as the difference of two continuous functions bounded below by 1, the result follows. □

C Isometries and submetries

5.6. Isometry. *Let $\mathcal{X}$ and $\mathcal{Y}$ be metric spaces and $\Phi\colon \mathcal{X} \to \mathcal{Y}$ be a map*

a) The map Φ is distance-preserving if

$$|\Phi(x) - \Phi(x')|_{\mathcal{Y}} = |x - x'|_{\mathcal{X}}$$

for any $x, x' \in X$.

b) A distance-preserving bijection Φ is called an isometry.

c) The spaces X and Y are called isometric (briefly $X \overset{iso}{=\!=} Y$) if there is an isometry $\Phi\colon X \to Y$.

5.7. Submetry. *A map* $\sigma\colon \mathcal{L}\to\mathcal{M}$ *between the metric spaces* $\mathcal{L}$ *and* $\mathcal{M}$ *is called a submetry if*

$$\sigma(\mathrm{B}(p,r)_{\mathcal{L}})=\mathrm{B}(\sigma(p),r)_{\mathcal{M}}$$

for any $p\in\mathcal{L}$ *and* $r\geqslant 0$.

Note $\sigma\colon \mathcal{L}\to\mathcal{M}$ is a submetry if it is 1-Lipschitz and 1-co-Lipschitz at the same time.

Note also that any submetry is an onto map.

The main source of examples of submetries comes from isometric group actions.

Namely, assume $\mathcal{L}$ is a metric space and G is a subgroup of isometries of $\mathcal{L}$. Denote by $\mathcal{L}/G$ the set of G-orbits; let us equip it with the pseudometric defined by

$$|G\cdot x-G\cdot y|_{\mathcal{L}/G}=\inf\left\{\,|g\cdot x-h\cdot y|_{\mathcal{L}}\,:\,g,h\in G\,\right\}.$$

Note that if all the G-orbits form closed sets in $\mathcal{L}$, then $\mathcal{L}/G$ is a genuine metric space.

5.8. Proposition. *Let* $\mathcal{L}$ *be a metric space. Assume that a group* G *acts on* $\mathcal{L}$ *by isometries and in such a way that every* G*-orbit is closed. Then the projection map* $\mathcal{L}\to\mathcal{L}/G$ *is a submetry.*

Proof. Denote by $\hat x$ the projection of $x\in\mathcal{L}$ in $\mathcal{L}/G$. We need to show that the map $x\mapsto\hat x$ is 1-Lipschitz and 1-co-Lipschitz. The co-Lipschitz part follows directly from the definitions of Hausdorff distance and co-Lipschitz maps.

Assume $|x-y|_{\mathcal{L}}<r$; equivalently $\mathrm{B}(x,r)_{\mathcal{L}}\ni y$. Since the action $G\curvearrowright\mathcal{L}$ is isometric, $\mathrm{B}(g\cdot x,r)_{\mathcal{L}}\ni g\cdot y$ for any $g\in G$.

In particular the orbit $G\cdot y$ lies in the open r-neighborhood of the orbit $G\cdot x$. In the same way we can prove that the orbit $G\cdot x$ lies in the open r-neighborhood of the orbit $G\cdot y$. That is, the Hausdorff distance between the orbits $G\cdot x$ and $G\cdot y$ is $<r$ or, equivalently, $|\hat x-\hat y|_{\mathcal{L}/G}<r$. Since x and y are arbitrary, the map $x\mapsto\hat x$ is 1-Lipschitz. □

5.9. Proposition. *Let* $\mathcal{L}$ *be a length space and* $\sigma\colon\mathcal{L}\to\mathcal{M}$ *be a submetry. Then* $\mathcal{M}$ *is a length space.*

Proof. Fix $\varepsilon>0$ and a pair of points $x,y\in\mathcal{M}$.

Since σ is 1-co-Lipschitz, there are points $\hat x,\hat y\in\mathcal{L}$ such that $\sigma(\hat x)=x$, $\sigma(\hat y)=y$, and $|\hat x-\hat y|_{\mathcal{L}}<|x-y|_{\mathcal{M}}+\varepsilon$.

Since $\mathcal{L}$ is a length space, there is a curve γ joining $\hat x$ to $\hat y$ in $\mathcal{L}$ such that

$$\operatorname{length}\gamma\leqslant|x-y|_{\mathcal{M}}+\varepsilon.$$

Since σ is 1-Lipschitz,

$$\operatorname{length}\sigma\circ\gamma \leqslant \operatorname{length}\gamma.$$

The curve $\sigma\circ\gamma$ joins x to y, and by the above,

$$\operatorname{length}\sigma\circ\gamma < |x-y|_{\mathcal{M}} + \varepsilon.$$

Since $\varepsilon > 0$ is arbitrary, $\mathcal{M}$ is a length space. □

D Speed of curves

Let $\mathcal{X}$ be a metric space. Recall that a curve in $\mathcal{X}$ is a continuous map $\alpha\colon \mathbb{I}\to\mathcal{X}$, where $\mathbb{I}$ is a real interval. A curve is called Lipschitz or locally Lipschitz if α is Lipschitz or locally Lipschitz. Length of curves is defined in 2.5.

The following theorem follows from [27, 2.7].

5.10. Theorem. *Let $\mathcal{X}$ be a metric space and $\alpha\colon\mathbb{I}\to\mathcal{X}$ be a locally Lipschitz curve. Then the speed function*

$$\operatorname{speed}_{t_0}\alpha = \lim_{\substack{t\to t_0+\\ s\to t_0-}} \frac{|\alpha(t)-\alpha(s)|}{|t-s|}$$

is defined for almost all $t_0\in\mathbb{I}$, and

$$\operatorname{length}\alpha = \int\limits_{\mathbb{I}} \operatorname{speed}_t\alpha\cdot \boldsymbol{d}t,$$

where $\int$ denotes the Lebesgue integral.

A curve $\alpha\colon\mathbb{I}\to\mathcal{X}$ is called a unit-speed curve if for any subinterval $[a,b]\subset\mathbb{I}$, we have

$$b-a = \operatorname{length}(\alpha|_{[a,b]}).$$

According to the above theorem, this is equivalent to the condition that α is Lipschitz and $\operatorname{speed}\alpha \overset{a.e.}{=\!=} 1$.

The following generalization of the standard Rademacher theorem on differentiability almost everywhere of Lipschitz maps between smooth manifolds [27, 5.5.2] was proved by Bernd Kirchheim [76].

The conclusion of the standard Rademacher theorem does not make sense for maps to a metric space since the target might have no linear structure. But the theorem does not hold even if we assume that the target is a Banach space. For example the map $[0,1]\to L^1[0,1]$ defined by $x\mapsto\chi_{[0,x]}$ is distance-preserving and in particular Lipschitz (here χ_A

denotes the characteristic function of A). However the differential $\boldsymbol{d}_x f$ is not defined for any x.

5.11. Theorem. *Let $\mathcal{X}$ be a metric space and $f\colon \mathbb{R}^n \hookrightarrow \mathcal{X}$ be 1-Lipschitz. Then for almost all $x \in \operatorname{Dom} f$ there is a pseudonorm $\|{*}\|_x$ on $\mathbb{R}^n$ such that*

$$|f(y) - f(z)|_{\mathcal{X}} = \|z - y\|_x + o(|y - x| + |z - x|).$$

Given f, the (pseudo)norm $\|{*}\|_x$ in the above theorem will be called its differential of the induced metric at x, or metric differential at x.

E Convex real-to-real functions

We will be interested in generalized solutions of the following differential inequalities

❶ $$y'' + \kappa{\cdot}y \geqslant \lambda \quad \text{and respectively} \quad y'' + \kappa{\cdot}y \leqslant \lambda$$

for fixed $\kappa, \lambda \in \mathbb{R}$. The solutions $y\colon \mathbb{R} \hookrightarrow \mathbb{R}$ are only assumed to be upper (respectively lower) semicontinuous subfunctions.

The inequalities ❶ are understood in the sense of distributions. That is, for any smooth function φ with compact support $\operatorname{Supp}\varphi \subset \operatorname{Dom} y$ the following inequality should be satisfied:

❷ $$\int\limits_{\operatorname{Dom} y} [y(t){\cdot}\varphi''(t) + \kappa{\cdot}y(t){\cdot}\varphi(t) - \lambda]\cdot \boldsymbol{d}t \geqslant 0$$
$$\text{respectively} \quad \leqslant 0.$$

The integral is understood in the sense of Lebesgue; in particular the inequality ❷ makes sense for any Borel-measurable subfunction y. The proofs of the following propositions are straightforward.

5.12. Proposition. *Let $\mathbb{I} \subset \mathbb{R}$ be an open interval and $y_n\colon \mathbb{I} \to \mathbb{R}$ be a sequence of solutions of one of the inequalities in ❶. Assume $y_n(t) \to y_\infty(t)$ as $n \to \infty$ for any $t \in \mathbb{I}$. Then y_∞ is a solution of the same inequality in ❶.*

Assume y is a solution of one of the inequalities in ❶. For $t_0 \in \operatorname{Dom} y$, let us define the right (left) derivative $y^+(t_0)$ ($y^-(t_0)$) at t_0 by

$$y^\pm(t_0) = \lim_{t \to t_0 \pm} \frac{y(t) - y(t_0)}{|t - t_0|}.$$

Note that our sign convention for y^- is not standard — for $y(t) = t$ we have $y^+(t) = 1$ and $y^-(t) = -1$.

5.13. Proposition. *Let $\mathbb{I} \subset \mathbb{R}$ be an open interval and $y\colon \mathbb{I} \to \mathbb{R}$ be a solution of an inequality in ❶. Then y is locally Lipschitz; its right and left derivatives $y^+(t_0)$ and $y^-(t_0)$ are defined for any $t_0 \in \mathbb{I}$. Moreover*

$$y^+(t_0) + y^-(t_0) \geqslant 0 \quad \textit{or respectively} \quad y^+(t_0) + y^-(t_0) \leqslant 0.$$

The next theorem gives a number of equivalent ways to define such generalized solutions.

5.14. Theorem. *Let $\mathbb{I}$ be an open real interval and $y\colon \mathbb{I} \to \mathbb{R}$ be a locally Lipschitz function. Then the following conditions are equivalent:*

a) $y'' \geqslant \lambda - \kappa{\cdot}y$ *(respectively* $y'' \leqslant \lambda - \kappa{\cdot}y$*).*

b) (barrier inequality) For any $t_0 \in \mathbb{I}$, there is a solution $\bar y$ of the ordinary differential equation $\bar y'' = \lambda - \kappa{\cdot}\bar y$ with $\bar y(t_0) = y(t_0)$ such that $\bar y \geqslant y$ (respectively $\bar y \leqslant y$) for all $t \in [t_0 - \varpi^\kappa, t_0 + \varpi^\kappa] \cap \mathbb{I}$.
The function $\bar y$ is called a lower (respectively upper) barrier of y at t_0.

c) (Jensen's inequality) For any pair of values $t_1 < t_2$ in $\mathbb{I}$ such that $|t_2 - t_1| < \varpi^\kappa$, the unique solution $z(t)$ of

$$z'' = \lambda - \kappa{\cdot}z$$

such that

$$z(t_1) = y(t_1), \quad z(t_2) = y(t_2)$$

satisfies $y(t) \leqslant z(t)$ (respectively $y(t) \geqslant z(t)$) for all $t \in [t_1, t_2]$.

Further, the following property holds:

d) Suppose $y'' \leqslant \lambda - \kappa{\cdot}y$. Let $t_0 \in \mathbb{I}$, and $\bar y$ be a solution of the ordinary differential equation $\bar y'' = \lambda - \kappa{\cdot}\bar y$ such that $\bar y(t_0) = y(t_0)$ and $y^+(t_0) \leqslant \bar y^+(t_0) \leqslant -y^-(t_0)$. (Note that such a $\bar y$ is unique if y is differentiable at t_0.)
Then $\bar y \geqslant y$ for all $t \in [t_0 - \varpi^\kappa, t_0 + \varpi^\kappa] \cap \mathbb{I}$; that is, $\bar y$ is a barrier of y at t_0. (Similarly, by reversing inequalities, for $y'' \geqslant \lambda - \kappa{\cdot}y$.)

The proof is left to the reader.

Note that Theorem 5.14 in particular implies that y satisfies $y'' \geqslant \lambda$ ($y'' \leqslant \lambda$) on an interval $\mathbb{I} \subset \mathbb{R}$ if and only if $y(t) - \frac{\lambda}{2}t^2$ is convex (concave) on $\mathbb{I}$.

We will often need the following fact about convergence of derivatives of convex functions:

5.15. Two-shoulder lemma. *Let $\mathbb{I}$ be an open interval and $f_n\colon \mathbb{I} \to \mathbb{R}$ be a sequence of convex functions. Assume the functions f_n pointwise* *converge to a function* $f\colon \mathbb{I}\to\mathbb{R}$*. Then for any* $t_0\in\mathbb{I}$,

$$f^{\pm}(t_0)\leqslant \varliminf_{n\to\infty} f_n^{\pm}(t_0).$$

Proof. Since the f_n are convex, we have $f_n^+(t_0)+f_n^-(t_0)\geqslant 0$, and for any t,

$$f_n(t)\geqslant f_n(t_0)\pm f^{\pm}(t_0)\cdot(t-t_0).$$

Passing to the limit, we get

$$f(t)\geqslant f(t_0)+\left[\varlimsup_{n\to\infty} f_n^+(t_0)\right]\cdot(t-t_0)$$

for $t\geqslant t_0$, and

$$f(t)\geqslant f(t_0)-\left[\varlimsup_{n\to\infty} f_n^-(t_0)\right]\cdot(t-t_0)$$

for $t\leqslant t_0$. Hence the result. □

5.16. Corollary. *Let* $\mathbb{I}$ *be an open interval and* $f_n\colon\mathbb{I}\to\mathbb{R}$ *be a sequence of functions such that* $f_n''\leqslant\lambda$ *that converge pointwise to a function* $f\colon\mathbb{I}\to\to\mathbb{R}$*. Then:*

a) If f *is differentiable at* $t_0\in\mathbb{I}$*, then*

$$f'(t_0)=\pm\lim_{n\to\infty} f_n^{\pm}(t_0).$$

b) If all f_n *and* f *are differentiable at* $t_0\in\mathbb{I}$*, then*

$$f'(t_0)=\lim_{n\to\infty} f_n'(t_0).$$

Proof. Set $\hat f_n(t)=f_n(t)-\frac{\lambda}{2}\cdot t^2$ and $\hat f(t)=f(t)-\frac{\lambda}{2}\cdot t^2$. Note that the $\hat f_n$ are concave and $\hat f_n\to\hat f$ pointwise. Thus the theorem follows from the two-shoulder lemma (5.15). □

F Convex functions on a metric space

In this section we define different types of convexity/concavity in the context of metric spaces; it will be mostly used for geodesic spaces. The notation refers to the corresponding second-order ordinary differential inequality.

5.17. Definition. *Let* $\mathcal{X}$ *be a metric space. We say that an upper semicontinuous subfunction* $f\colon\mathcal{X}\hookrightarrow(-\infty,\infty]$ *satisfies the inequality*

$$f''+\kappa\cdot f\geqslant\lambda$$

if for any unit-speed geodesic γ in $\operatorname{Dom} f$, *the real-to-real function* $y(t) = f \circ \gamma(t)$ *satisfies*

$$y'' + \kappa \cdot y \geqslant \lambda$$

in the domain $\{\, t : y(t) < \infty \,\}$*; see the definition in Section 5E.*

We say that a lower semicontinuous subfunction $f\colon \mathcal{X} \hookrightarrow [-\infty, \infty)$ *satisfies the inequality*

$$f'' + \kappa \cdot f \leqslant \lambda$$

if the subfunction $h = -f$ *satisfies*

$$h'' - \kappa \cdot h \geqslant -\lambda.$$

Functions satisfying the inequalities

$$f'' \geqslant \lambda \quad \text{and} \quad f'' \leqslant \lambda$$

are called λ*-convex and* λ*-concave respectively.*

0*-convex and* 0*-concave subfunctions will also be called convex and concave respectively.*

If f *is* λ*-convex for some* $\lambda > 0$, f *will be called strongly convex; correspondingly, if* f *is* λ*-concave for some* $\lambda < 0$, f *will be called strongly concave.*

If for any point $p \in \operatorname{Dom} f$ *there is a neighborhood* $\Omega \ni p$ *and a real number* λ *such that the restriction* $f|_{\Omega}$ *is* λ*-convex (or* λ*-concave), then* f *is called semiconvex (respectively semiconcave).*

Various authors define the class of λ-convex (λ-concave) functions differently. Their definitions may correspond to $\pm\lambda$-convex ($\pm\lambda$-concave) or $\pm\frac{\lambda}{2}$-convex ($\pm\frac{\lambda}{2}$-concave) functions in our definitions.

5.18. Proposition. *Let* $\mathcal{X}$ *be a metric space. Assume that* $f\colon \mathcal{X} \hookrightarrow \mathbb{R}$ *is a semiconvex subfunction and* $\varphi\colon \mathbb{R} \to \mathbb{R}$ *is a nondecreasing semiconvex function. Then the composition* $\varphi \circ f$ *is a semiconvex subfunction.*

The proof is straightforward.

Chapter 6

The ghost of Euclid

A Geodesics, triangles and hinges

Geodesics and their relatives. Let $\mathcal{X}$ be a metric space and $\mathbb{I} \subset \mathbb{R}$ be an interval. A globally distance-preserving map $\gamma\colon \mathbb{I} \to \mathcal{X}$ is called a unit-speed geodesic[1]; in other words, $\gamma\colon \mathbb{I} \to \mathcal{X}$ is a unit-speed geodesic if the equality

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

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

A unit-speed geodesic $\gamma\colon \mathbb{R}_{\geqslant 0} \to \mathcal{X}$ is called a half-line.

A unit-speed geodesic $\gamma\colon \mathbb{R} \to \mathcal{X}$ is called a line.

A unit-speed geodesic between p and q in $\mathcal{X}$ will be denoted by $\mathrm{geod}_{[pq]}$. We will always assume $\mathrm{geod}_{[pq]}$ is parametrized starting at p; that is, $\mathrm{geod}_{[pq]}(0) = p$ and $\mathrm{geod}_{[pq]}(|p - q|) = q$. The image of $\mathrm{geod}_{[pq]}$ will be denoted by $[pq]$ and called a geodesic. The term geodesic will also be used for a linear reparametrization of a unit-speed geodesic; when confusion is possible we will call the latter a pregeodesic. With a slight abuse of notation, we will use the notation $[pq]$ also for the class of all linear reparametrizations of $\mathrm{geod}_{[pq]}$.

We may write $[pq]_{\mathcal{X}}$ to emphasize that the geodesic $[pq]$ is in the space $\mathcal{X}$. Also we use the following short-cut notation:

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

In general, a geodesic between p and q need not exist and if it exists, it need not to be unique. However, once we write $\mathrm{geod}_{[pq]}$ or $[pq]$ we mean that we have fixed a choice of geodesic.

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

A constant-speed geodesic $\gamma\colon[0,1]\to\mathcal{X}$ is called a geodesic path. Given a geodesic $[pq]$, we denote by $\mathrm{path}_{[pq]}$ the corresponding geodesic path; that is,

$$\mathrm{path}_{[pq]}(t)\equiv\mathrm{geod}_{[pq]}(t{\cdot}|p-q|).$$

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 $U\ni t$ in $\mathbb{I}$ such that the restriction $\gamma|_U$ is a constant-speed geodesic. If $\mathbb{I}=[0,1]$, then γ is called a local geodesic path.

6.1. Proposition. *Suppose $\mathcal{X}$ is a metric space and $\gamma\colon[0,\infty)\to\mathcal{X}$ is a half-line. Then the Busemann function $\mathrm{bus}_\gamma\colon\mathcal{X}\to\mathbb{R}$*

❶ $$\mathrm{bus}_\gamma(x)=\lim_{t\to\infty}|\gamma(t)-x|-t$$

is defined and 1-Lipschitz.

Proof. As follows from the triangle inequality, the function

$$t\mapsto|\gamma(t)-x|-t$$

is nonincreasing in t. Clearly $|\gamma(t)-x|-t\geqslant-|\gamma(0)-x|$. Thus the limit in ❶ is defined. □

Triangles. 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, and 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. Or there may be many different triangles, any of which can be denoted by $[pqr]$. Once we write $[pqr]$, it means we have chosen such a triangle; that is, made a choice of each $[qr]$, $[rp]$, and $[pq]$.

The value

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

will be called the perimeter of triangle $[pqr]$; it obviously coincides with perimeter of the triple p, q, r as defined below.

Hinges. 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 $[p\,{}^x_y]=([px],[py])$.

B Model angles and triangles

Let $\mathcal{X}$ be a metric space, $p,q,r\in\mathcal{X}$, and $\kappa\in\mathbb{R}$. Let us define the model triangle $[\tilde p\tilde q\tilde r]$ (briefly, $[\tilde p\tilde q\tilde r]=\tilde\triangle^\kappa(pqr)$) to be a triangle in the model plane $\mathbb{M}^2(\kappa)$ such that

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

In the notation of Section 1A, $\tilde\triangle^\kappa(pqr) = \tilde\triangle^\kappa\{|q-r|, |r-p|, |p-q|\}$.

If $\kappa \leqslant 0$, the model triangle is always defined, that is, it exists and is unique up to isometry of $\mathbb{M}^2(\kappa)$. If $\kappa > 0$, the model triangle is said to be defined if in addition

$$|p-q| + |q-r| + |r-p| < 2\cdot\varpi^\kappa.$$

In this case, the model triangle exists and is unique up to isometry of $\mathbb{M}^2(\kappa)$. The value $|p-q| + |q-r| + |r-p|$ will be called the perimeter of the triple p, q, r.

If for $p, q, r \in \mathcal{X}$, $[\tilde p\tilde q\tilde r] = \tilde\triangle^\kappa(pqr)$ is defined 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 $\tilde\measuredangle^\kappa(p\,{}^q_r)$.

In the notation of Section 1A,

$$\tilde\measuredangle^\kappa(p\,{}^q_r) = \tilde\measuredangle^\kappa\{|q-r|; |p-q|, |p-r|\}.$$

6.2. Alexandrov's lemma. *Let p, q, r, z be distinct points in a metric space such that $z \in]pr[$ and*

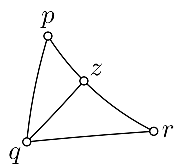

$$|p-q| + |q-r| + |r-p| < 2\cdot\varpi^\kappa.$$

Then the following expressions have the same sign:

a) $\tilde\measuredangle^\kappa(p\,{}^q_r) - \tilde\measuredangle^\kappa(p\,{}^q_z)$,

b) $\tilde\measuredangle^\kappa(z\,{}^q_p) + \tilde\measuredangle^\kappa(z\,{}^q_r) - \pi$.

Moreover,

$$\tilde\measuredangle^\kappa(q\,{}^p_r) \geqslant \tilde\measuredangle^\kappa(q\,{}^p_z) + \tilde\measuredangle^\kappa(q\,{}^z_r),$$

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

Proof. By the triangle inequality,

$$|p-q| + |q-z| + |z-p| \leqslant |p-q| + |q-r| + |r-p| < 2\cdot\varpi^\kappa.$$

Therefore the model triangle $[\tilde p\tilde q\tilde z] = \tilde\triangle^\kappa pqz$ is defined. Take a point $\tilde r$ on the extension of $[\tilde p\tilde z]$ beyond $\tilde z$ so that $|\tilde p-\tilde r| = |p-r|$ (and therefore $|\tilde p-\tilde z| = |p-z|$).

From monotonicity of the function $a \mapsto \tilde\measuredangle^\kappa\{a; b, c\}$ (1.1*c*), the following expressions have the same sign:

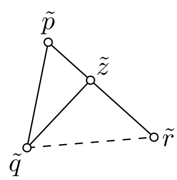

(i) $\measuredangle[\tilde p\,{}^{\tilde q}_{\tilde r}] - \tilde\measuredangle^\kappa(p\,{}^q_r)$;

(ii) $|\tilde p-\tilde r| - |p-r|$;

(iii) $\measuredangle[\tilde z\,{}^{\tilde q}_{\tilde r}] - \tilde\measuredangle^\kappa(z\,{}^q_r)$.

Since

$$\measuredangle[\tilde{p}\,{}^{\tilde{q}}_{\tilde{r}}] = \measuredangle[\tilde{p}\,{}^{\tilde{q}}_{\tilde{z}}] = \tilde{\measuredangle}^{\kappa}(p\,{}^{q}_{z})$$

and

$$\measuredangle[\tilde{z}\,{}^{\tilde{q}}_{\tilde{r}}] = \pi - \measuredangle[\tilde{z}\,{}^{\tilde{p}}_{\tilde{q}}] = \pi - \tilde{\measuredangle}^{\kappa}(z\,{}^{p}_{q}),$$

the first statement follows.

For the second statement, construct $[\tilde{q}\tilde{z}r'] = \tilde{\triangle}^{\kappa}qzr$ on the opposite side of $[\tilde{q}\tilde{z}]$ from $[\tilde{p}\tilde{q}\tilde{z}]$. Since

$$|\tilde{p} - r'| \leqslant |\tilde{p} - \tilde{z}| + |\tilde{z} - r'| = |p - z| + |z - r| = |p - r|,$$

then

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

Equality holds if and only if $|\tilde{p} - r'| = |p - r|$, as required. □

C Angles and the first variation

Given a hinge $[p\,{}^{x}_{y}]$, we define its *angle* to be

❶ $$\measuredangle[p\,{}^{x}_{y}] := \lim_{\bar{x},\bar{y}\to p} \tilde{\measuredangle}^{\kappa}(p\,{}^{\bar{x}}_{\bar{y}}),$$

for $\bar{x} \in\,]px]$ and $\bar{y} \in\,]py]$, if this limit exists.

Similarly to $\tilde{\measuredangle}^{\kappa}(p\,{}^{x}_{y})$, we will use the short notation

$$\tilde{\curlyvee}^{\kappa}[p\,{}^{x}_{y}] = \tilde{\curlyvee}^{\kappa}\left\{\measuredangle[p\,{}^{x}_{y}]; |p - x|, |p - y|\right\},$$

where the right-hand side is defined in Section 1A. The value $\tilde{\curlyvee}^{\kappa}[p\,{}^{x}_{y}]$ will be called the *model side* of the hinge $[p\,{}^{x}_{y}]$.

6.3. Lemma. *Let p, x, y be a triple of points in a metric space with perimeter ℓ. Then for any $\kappa, \mathrm{K} \in \mathbb{R}$,*

❷, $$|\tilde{\measuredangle}^{\mathrm{K}}(p\,{}^{x}_{y}) - \tilde{\measuredangle}^{\kappa}(p\,{}^{x}_{y})| \leqslant 100(|\mathrm{K}| + |\kappa|)\cdot\ell^2,$$

whenever the left-hand side is defined.

Lemma 6.3 implies that the definition of angle is independent of κ. In particular, one can take $\kappa = 0$ in ❶; thus the angle can be calculated from the cosine law:

$$\cos\tilde{\measuredangle}^{0}(p\,{}^{x}_{y}) = \frac{|p - x|^2 + |p - y|^2 - |x - y|^2}{2\cdot|p - x|\cdot|p - y|}.$$

Proof. The function $\kappa \mapsto \tilde{\measuredangle}^\kappa(p\,{}^x_y)$ is nondecreasing (1.1*d*). Thus, for $K > \kappa$, we have

$$0 \leqslant \tilde{\measuredangle}^K(p\,{}^x_y) - \tilde{\measuredangle}^\kappa(p\,{}^x_y) \leqslant \tilde{\measuredangle}^K(p\,{}^x_y) + \tilde{\measuredangle}^K(x\,{}^p_y) + \tilde{\measuredangle}^K(y\,{}^p_x) - \\ - \tilde{\measuredangle}^\kappa(p\,{}^x_y) - \tilde{\measuredangle}^\kappa(x\,{}^p_y) - \tilde{\measuredangle}^\kappa(y\,{}^p_x) = \\ = K\cdot \operatorname{area} \tilde{\triangle}^K(pxy) - \kappa\cdot \operatorname{area} \tilde{\triangle}^\kappa(pxy).$$

Note that for $\kappa \geqslant 0$ a triangle of perimeter l in $\mathbb{M}^2(\kappa)$ lies in a ball of radius $2l$, which easily implies that $\operatorname{area} \tilde{\triangle}^\kappa(pxy) \leqslant 100\ell^2$. For $\kappa < 0$ one gets the same estimate by a direct computation in the hyperbolic plane.

Therefore

$$\operatorname{area} \tilde{\triangle}^\kappa(pxy) \leqslant 100\ell^2, \qquad \operatorname{area} \tilde{\triangle}^K(pxy) \leqslant 100\ell^2.$$

Thus ❷ follows. □

6.4. Triangle inequality for angles. *Let $[px^1]$, $[px^2]$, and $[px^3]$ be three geodesics in a metric space. If all of 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 can assume that $\alpha^{12} + \alpha^{23} < \pi$. Set $\gamma^i = \operatorname{geod}_{[px^i]}$. Given any $\varepsilon > 0$, for all sufficiently small $t, \tau, s \in \mathbb{R}_+$ we have

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

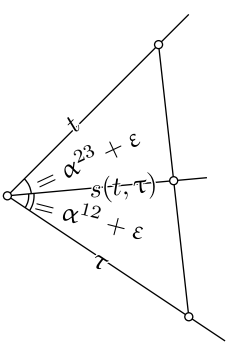

Below we define $s(t,\tau)$ so that for $s = s(t,\tau)$, this chain of inequalities can be continued as follows:

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

Thus for any $\varepsilon > 0$,

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

Hence the result follows.

To define $s(t,\tau)$, consider three half-lines $\tilde\gamma^1$, $\tilde\gamma^2$, $\tilde\gamma^3$ 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 half-line 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. □

6.5. 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.

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

$$|p-\mathrm{geod}_{[qx]}(t)|\leqslant|q-p|-t\cdot\cos\alpha+o(t).$$

Proof. Take a sufficiently small $\varepsilon>0$. For all sufficiently small $t>0$, we have

$$\begin{aligned}|\mathrm{geod}_{[qp]}(t/\varepsilon)-\mathrm{geod}_{[qx]}(t)|&\leqslant\tfrac{t}{\varepsilon}\cdot\sqrt{1+\varepsilon^2-2\cdot\varepsilon\cdot\cos\alpha}+o(t)\leqslant\\&\leqslant\tfrac{t}{\varepsilon}-t\cdot\cos\alpha+t\cdot\varepsilon.\end{aligned}$$

Applying the triangle inequality, we get

$$\begin{aligned}|p-\mathrm{geod}_{[qx]}(t)|&\leqslant|p-\mathrm{geod}_{[qp]}(t/\varepsilon)|+|\mathrm{geod}_{[qp]}(t/\varepsilon)-\mathrm{geod}_{[qx]}(t)|\leqslant\\&\leqslant|p-q|-t\cdot\cos\alpha+t\cdot\varepsilon\end{aligned}$$

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

D Space of directions

Let $\mathcal{X}$ be a metric space. If the angle $\measuredangle[p\,{}^x_y]$ is defined for any hinge $[p\,{}^x_y]$ in $\mathcal{X}$, then we will say that the space $\mathcal{X}$ has defined angles.

Let $\mathcal{X}$ be a space with defined angles. For $p\in\mathcal{X}$, consider the set $\mathfrak{S}_p$ of all nontrivial unit-speed geodesics starting at p. By 6.4, the triangle inequality holds for $\measuredangle$ on $\mathfrak{S}_p$, that is, $(\mathfrak{S}_p,\measuredangle)$ forms a pseudometric space.

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

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

the direction of $[px]$ is denoted by $\uparrow_{[px]}$.

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

E Tangent space

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

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

where $\vartheta = \min\{\pi, |p - q|_{\mathcal{X}}\}$. We write $(t,x), t \neq 0$, as $t{\cdot}x$, referred to as cone multiplication.

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

The scalar product $\langle v, w\rangle$ of two vectors $v = (s,p)$ and $w = (t,q)$ in $\operatorname{Cone}\mathcal{X}$ is defined by

$$\langle v, w\rangle := |v|{\cdot}|w|{\cdot}\cos\vartheta;$$

we set $\langle v, w\rangle := 0$ if $v = 0$ or $w = 0$.

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

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

$$\text{❶.}\qquad |\alpha - \beta|_{\mathfrak{T}_p} = \lim_{\varepsilon\to 0} \frac{|\alpha(\varepsilon) - \beta(\varepsilon)|_{\mathcal{X}}}{\varepsilon}.$$

If the angles in $\mathcal{X}$ are defined, then so is the limit in ❶, and we obtain 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 formed by the equivalence classes for the relation

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

The completion of T'_p is therefore naturally isometric to T_p. The element of T'_p that corresponds to the geodesic path $\operatorname{geod}_{[pq]}$ is called logarithm of $[pq]$ and denoted by $\log[pq]$.

The elements of T_p will be called tangent vectors at p, despite that T_p is only a cone — not a vector space. The elements of T'_p will be called geodesic tangent vectors at p. The tip of the tangent cone T_p will be denoted by 0 or 0_p.

F Velocity of curves

6.7. Definition. *Let $\mathcal{X}$ be a metric space. Let $\alpha\colon [0,a)\to\mathcal{X}$ for some $a>0$ be a function, not necessarily continuous, such that $\alpha(0)=p$. We say that $v\in\mathrm{T}_p$ is the right derivative of α at 0, briefly $\alpha^+(0)=v$, if for some (and therefore any) sequence of vectors $v_n\in\mathrm{T}'_p$ such that $v_n\to v$ as $n\to\infty$, and corresponding geodesics γ_n, we have*

$$\varlimsup_{\varepsilon\to 0+}\frac{|\alpha(\varepsilon)-\gamma_n(\varepsilon)|_{\mathcal{X}}}{\varepsilon}\to 0\quad\text{as}\quad n\to\infty.$$

We define right and left derivatives $\alpha^+(t_0)$ and $\alpha^-(t_0)$ of α at $t_0\in\mathbb{I}$ by

$$\alpha^{\pm}(t_0)=\check{\alpha}^+(0),$$

where $\check{\alpha}(t)=\alpha(t_0\pm t)$.

The sign convention is not quite standard; if α is a smooth curve in a Riemannian manifold then we have

$$\alpha^+(t)=-\alpha^-(t).$$

Note that if γ is a geodesic starting at p and the tangent vector $v\in\mathrm{T}'_p$ corresponds to γ, then $\gamma^+(0)=v$.

6.8. Exercise. *Assume $\mathcal{X}$ is a metric space with defined angles, and let $\alpha,\beta\colon[0,a)\to\mathcal{X}$ be two maps such that the right derivatives $\alpha^+(0)$, $\beta^+(0)$ are defined and $\alpha^+(0)=\beta^+(0)$. Show that*

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

6.9. Proposition. *Let $\mathcal{X}$ be a metric space with defined angles and $p\in\in\mathcal{X}$. Then for any tangent vector $v\in\mathrm{T}_p\mathcal{X}$ there is a map $\alpha\colon[0,\varepsilon)\to\mathcal{X}$ such that $\alpha^+(0)=v$.*

Proof. If $v\in\mathrm{T}'_p$, then for the corresponding geodesic α we have $\alpha^+(0)=\,=v$.

Given $v\in\mathrm{T}_p$, construct a sequence $v_n\in\mathrm{T}'_p$ such that $v_n\to v$, and let γ_n be a sequence of corresponding geodesics.

The needed map α can be found among the maps such that $\alpha(0) = p$ and

$$\alpha(t) = \gamma_n(t) \quad \text{if} \quad \varepsilon_{n+1} \leqslant t < \varepsilon_n,$$

where ε_n is a decreasing sequence converging to 0 as $n \to \infty$. In order to satisfy the conclusion of the proposition, one has to choose the sequence ε_n converging to 0 very fast. □

6.10. Definition. *Let $\mathcal{X}$ be a metric space and $\alpha\colon \mathbb{I} \to \mathcal{X}$ be a curve.*

For $t_0 \in \mathbb{I}$, if $\alpha^+(t_0)$ or $\alpha^-(t_0)$ or both are defined, we say respectively that α is right or left or both-sided differentiable at t_0. In the exceptional cases where t_0 is the left (respectively right) end of $\mathbb{I}$, α is by definition left (respectively right) differentiable at t_0.

If α is both-sided differentiable at t, and

$$|\alpha^+(t)| = |\alpha^-(t)| = \tfrac{1}{2}\cdot|\alpha^+(t) - \alpha^-(t)|_{\mathrm{T}_{\alpha(t)}},$$

then we say that α is differentiable at t.

6.11. Exercise. *Assume $\mathcal{X}$ is a metric space with defined angles. Show that any geodesic $\gamma\colon \mathbb{I} \to \mathcal{X}$ is differentiable everywhere.*

Recall that the speed of a curve is defined in 5.10.

6.12. Exercise. *Let α be a curve in a metric space with defined angles. Suppose that $\operatorname{speed}_t \alpha$, $\alpha^+(t)$, and $\alpha^-(t)$ are defined.*

Show that α is differentiable at t.

G Differential

6.13. Definition. *Let $\mathcal{X}$ be a metric space with defined angles, and $f\colon \mathcal{X} \hookrightarrow \mathbb{R}$ be a subfunction. For $p \in \operatorname{Dom} f$, a function $\varphi\colon \mathrm{T}_p \to \mathbb{R}$ is called the differential of f at p (briefly $\varphi = \boldsymbol{d}_p f$) if for any map $\alpha\colon \mathbb{I} \to \mathcal{X}$ such that $\mathbb{I}$ is a real interval, $\alpha(0) = p$, and $\alpha^+(0)$ is defined, we have*

$$(f \circ \alpha)^+(0) = \varphi(\alpha^+(0)).$$

6.14. Proposition. *Let $f\colon \mathcal{X} \hookrightarrow \mathbb{R}$ be a locally Lipschitz semiconcave subfunction on a metric space $\mathcal{X}$ with defined angles. Then the differential $\boldsymbol{d}_p f$ is uniquely defined for any $p \in \operatorname{Dom} f$. Moreover,*

a) the differential $\boldsymbol{d}_p f\colon \mathrm{T}_p \to \mathbb{R}$ is Lipschitz and the Lipschitz constant of $\boldsymbol{d}_p f\colon \mathrm{T}_p \to \mathbb{R}$ does not exceed the Lipschitz constant of f in a neighborhood of p.

b) $\boldsymbol{d}_p f\colon \mathrm{T}_p \to \mathbb{R}$ *is a positive homogeneous function; that is, for any* $\mathfrak{z} \geqslant 0$ *and* $v \in \mathrm{T}_p$ *we have*

$$\mathfrak{z}\cdot\boldsymbol{d}_p f(v) = \boldsymbol{d}_p f(\mathfrak{z}\cdot v).$$

c)

$$\boldsymbol{d}_p f = \boldsymbol{d}_p^{\omega} f|_{\mathrm{T}_p}.$$

Proof. Passing to a subdomain of f if necessary, we can assume that f is ℓ-Lipschitz and λ-concave for some $\ell, \lambda \in \mathbb{R}$.

Take a geodesic γ in $\operatorname{Dom} f$ starting at p. Since $f\circ\gamma$ is semiconcave, the right derivative $(f\circ\gamma)^+(0)$ is defined. Since f is ℓ-Lipschitz, we have

❶ $$|(f\circ\gamma)^+(0) - (f\circ\gamma_1)^+(0)| \leqslant \ell\cdot|\gamma^+(0) - \gamma_1^+(0)|$$

for any other geodesic γ_1 starting at p.

Define $\varphi\colon \mathrm{T}'_p \to \mathbb{R}\colon \gamma^+(0) \mapsto (f\circ\gamma)^+(0)$. From ❶, φ is an ℓ-Lipschtz function defined on T'_p. Thus we can extend φ to all of T_p as an ℓ-Lipschitz function.

It remains to show that φ is the differential of f at p. Assume $\alpha\colon [0,a) \to \mathcal{X}$ is a map such that $\alpha(0) = p$ and $\alpha^+(0) = v \in \mathrm{T}_p$. Let $\gamma_n \in \Gamma_p$ be a sequence of geodesics as in the definition 6.7; that is, if

$$v_n = \gamma_n^+(0) \quad \text{and} \quad a_n = \overline{\lim_{t\to 0+}}\, |\alpha(t) - \gamma_n(t)|/t$$

then $a_n \to 0$ and $v_n \to v$ as $n \to \infty$. Then

$$\varphi(v) = \lim_{n\to\infty} \varphi(v_n),$$

$$f\circ\gamma_n(t) = f(p) + \varphi(v_n)\cdot t + o(t),$$
$$|f\circ\alpha(t) - f\circ\gamma_n(t)| \leqslant \ell\cdot|\alpha(t) - \gamma_n(t)|.$$

Hence

$$f\circ\alpha(t) = f(p) + \varphi(v)\cdot t + o(t)$$

The last part follows from the definitions of differential and ultradifferential; see Section 3E. ☐

H Ultratangent space

Fix a selective ultrafilter ω on the set $\mathbb{N}$ of natural numbers.

For a metric space $\mathcal{X}$ and a positive real number $\mathfrak{z}$, we will denote by $\mathfrak{z}\cdot\mathcal{X}$ its $\mathfrak{z}$-b l o w u p, which is a metric space with the same underlying set as $\mathcal{X}$ and the metric multiplied by $\mathfrak{z}$. The tautological bijection $\mathcal{X} \to \mathfrak{z}\cdot\mathcal{X}$ will be denoted by $x \mapsto x^{\mathfrak{z}}$, so

$$|x^{\mathfrak{z}} - y^{\mathfrak{z}}| = \mathfrak{z}\cdot|x - y|$$

for any $x, y \in \mathcal{X}$.

The ϖ-blowup $\varpi\cdot\mathcal{X}$ of $\mathcal{X}$ is defined to be the ϖ-limit of the n-blowups $n\cdot\mathcal{X}$; that is,

$$\varpi\cdot\mathcal{X} := \lim_{n\to\varpi} n\cdot\mathcal{X}.$$

Given a point $x \in \mathcal{X}$, we can consider the sequence x^n, where $x^n \in$ $\in n\cdot\mathcal{X}$ is the image of x under n-blowup. Note that if $x \neq y$, then

$$|x^{\varpi} - y^{\varpi}|_{\varpi\cdot\mathcal{X}} = \infty;$$

that is, x^{ϖ} and y^{ϖ} belong to different metric components of $\varpi\cdot\mathcal{X}$.

The metric component of x^{ϖ} in $\varpi\cdot\mathcal{X}$ is called the ultratangent space of $\mathcal{X}$ at x and is denoted by $\mathrm{T}^{\varpi}_x\mathcal{X}$ or T^{ϖ}_x.

Equivalently, the ultratangent space $\mathrm{T}^{\varpi}_x\mathcal{X}$ can be defined as follows. Consider all the sequences of points $x_n \in \mathcal{X}$ such that the sequence $(n\cdot|x - x_n|_{\mathcal{X}})$ is bounded. Define the pseudodistance between two such sequences as

$$|(x_n) - (y_n)| = \lim_{n\to\varpi} n\cdot|x_n - y_n|_{\mathcal{X}}.$$

Then $\mathrm{T}^{\varpi}_x\mathcal{X}$ is the corresponding metric space.

Tangent spaces (see section 6E) as well as ultratangent spaces generalize the notion of tangent spaces on Riemannian manifolds. In the simplest cases these two notions define the same space. However in general they are different and are both useful — often a lack of some property in one is compensated by the other.

It is clear from the definition that a tangent space has a cone structure. On the other hand, in general an ultratangent space does not have a cone structure. Hilbert's cube $\prod_{n=1}^{\infty}[0, 2^{-n}]$ is an example. We remark that Hilbert's cube is a CBB(0) as well as a CAT(0) Alexandrov space.

The next theorem shows that the tangent space T_p can be (and often will be) considered as a subset of T^{ϖ}_p.

6.15. Theorem. *Let $\mathcal{X}$ be a metric space with defined angles. Then for any $p \in \mathcal{L}$, there is a distance-preserving map*

$$\iota : \mathrm{T}_p \hookrightarrow \mathrm{T}^{\varpi}_p$$

such that for any geodesic γ starting at p we have

$$\gamma^+(0) \overset{\iota}{\mapsto} \lim_{n\to\varpi}[\gamma(\tfrac{1}{n})]^n.$$

Proof. Given $v \in \mathrm{T}'_p$, choose a geodesic γ that starts at p and such that $\gamma^+(0) = v$. Set $v^n = [\gamma(\tfrac{1}{n})]^n \in n{\cdot}\mathcal{X}$ and

$$v^\omega = \lim_{n\to\omega} v^n.$$

Note that the value $v^\omega \in \mathrm{T}^\omega_p$ does not depend on the choice of γ; that is, if γ_1 is another geodesic starting at p such that $\gamma_1^+(0) = v$, then

$$\lim_{n\to\omega} v^n = \lim_{n\to\omega} v_1^n,$$

where $v_1^n = [\gamma_1(\tfrac{1}{n})]^n \in n{\cdot}\mathcal{X}$. The latter follows since

$$|\gamma(t) - \gamma_1(t)|_{\mathcal{X}} = o(t),$$

and therefore $|v^n - v_1^n|_{n\cdot\mathcal{X}} \to 0$ as $n \to \infty$.

Set $\iota(v) = v^\omega$. Since angles between geodesics in $\mathcal{X}$ are defined, for any $v, w \in \mathrm{T}'_p$ we have $n{\cdot}|v_n - w_n| \to |v - w|$. Thus $|v_\omega - w_\omega| = |v - w|$; that is, $\iota\colon \mathrm{T}'_p \to \mathrm{T}_p$ is a distance-preserving map.

Since T'_p is dense in T_p, we can extend ι to a distance-preserving map $\mathrm{T}_p \to \mathrm{T}^\omega_p$. □

I Ultradifferential

Given a function $f\colon \mathcal{L} \to \mathbb{R}$, consider the sequence of functions $f_n\colon n{\cdot}\mathcal{L} \to \to \mathbb{R}$ defined by

$$f_n(x^n) = n{\cdot}(f(x) - f(p)),$$

where $x^n \in n{\cdot}\mathcal{L}$ is the point correspnding to $x \in \mathcal{L}$. While $n{\cdot}(\mathcal{L}, p) \to \to (\mathrm{T}^\omega, 0)$ as $n \to \omega$, the functions f_n converge to the ω-differential of f at p. It will be denoted by $\boldsymbol{d}^\omega_p f$:

$$\boldsymbol{d}^\omega_p f\colon \mathrm{T}^\omega_p \to \mathbb{R}, \quad \boldsymbol{d}^\omega_p f = \lim_{n\to\omega} f_n.$$

Clearly, the ω-differential $\boldsymbol{d}^\omega_p f$ of a locally Lipschitz subfunction f is defined and Lipschitz at each point $p \in \operatorname{Dom} f$.

J Remarks

In Alexandrov geometry, angles defined as in Section 6C always exist (see Theorem 8.14*c* and Corollary 9.15*b*).

For general metric spaces, angles may not exist, and given a hinge $[p\,^x_y]$ it is more natural to consider the upper angle defined by

$$\measuredangle^{\mathrm{up}}[p\,^x_y] := \varlimsup_{\bar{x},\bar{y}\to p} \tilde{\measuredangle}^\kappa(p\,^{\bar{x}}_{\bar{y}}),$$

where $\bar{x} \in\,]px]$ and $\bar{y} \in\,]py]$. The triangle inequality (6.4) holds for upper angles as well.

Chapter 7

Dimension theory

In Section 7A, we give definitions of different types of dimension-like invariants of metric spaces and state general relations between them.

A Definitions

The proofs of most of the statements in this section can be found in the book of Witold Hurewicz and Henry Wallman [69]; the rest follow directly from the definitions.

7.1. Hausdorff dimension. *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.*

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 nonempy 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 $W_\alpha\subset$ $\subset V_\beta$.

7.2. Topological dimension. *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 finite 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 7.3 and 7.4 are commonly called “dimension”; for that reason we call these statements axioms.

7.3. Normalization axiom. *For any* $m \in \mathbb{Z}_{\geqslant 0}$,

$$\operatorname{TopDim}\mathbb{E}^m = \operatorname{HausDim}\mathbb{E}^m = m.$$

7.4. Cover axiom. *If* $\{A_n\}_{n=1}^{\infty}$ *is a countable closed cover of* $\mathcal{X}$, *then*

$$\begin{aligned}\operatorname{TopDim}\mathcal{X} &= \sup_n\{\operatorname{TopDim}A_n\},\\ \operatorname{HausDim}\mathcal{X} &= \sup_n\{\operatorname{HausDim}A_n\}.\end{aligned}$$

On product spaces. Recall that the direct product $\mathcal{X} \times \mathcal{Y}$ of metric spaces $\mathcal{X}$ and $\mathcal{Y}$ is defined in Section 11C. Direct product satisfies the following two inequalities:

$$\operatorname{TopDim}(\mathcal{X} \times \mathcal{Y}) \leqslant \operatorname{TopDim}\mathcal{X} + \operatorname{TopDim}\mathcal{Y}$$

and

$$\operatorname{HausDim}(\mathcal{X} \times \mathcal{Y}) \geqslant \operatorname{HausDim}\mathcal{X} + \operatorname{HausDim}\mathcal{Y}.$$

These inequalities might be strict. For topological dimension, strict inequality holds for a pair of Pontryagin surfaces [113]. For Hausdorff dimension, an example was constructed by Abram Besicovitch and Pat Moran [20].

The following theorem follows from [69, theorems V 8 and VII 2].

7.5. 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, $\operatorname{TopDim}\mathcal{X} \leqslant \operatorname{HausDim}\mathcal{X}$.

In fact it is true that for any separable metric space $\mathcal{X}$ we have

$$\operatorname{TopDim}\mathcal{X} = \inf\{\operatorname{HausDim}\mathcal{Y}\},$$

where the infimum is taken over all metric spaces $\mathcal{Y}$ homeomorphic to $\mathcal{X}$.

7.6. 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 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 [69, theorems VI 1&2]. (This theorem also holds for non-separable metric spaces [97], [49, 3.2.10]).

7.7. 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.*

7.8. Proposition. *Suppose $\mathcal{X}$ and $\mathcal{Y}$ are metric spaces and $\Phi\colon \mathcal{X} \to \mathcal{Y}$ satisfies*

$$|\Phi(x) - \Phi(x')| \geqslant \varepsilon\cdot|x - x'|$$

for some fixed $\varepsilon > 0$ and any pair $x, x' \in \mathcal{X}$. Then

$$\operatorname{HausDim}\mathcal{X} \leqslant \operatorname{HausDim}\mathcal{Y}.$$

In particular, if there is a Lipschitz onto map $\mathcal{Y} \to \mathcal{X}$, then

$$\operatorname{HausDim}\mathcal{X} \leqslant \operatorname{HausDim}\mathcal{Y}.$$

B Linear dimension

In addition to HausDim and TopDim, we will use yet another dimension, which we call linear dimension. It will be applied only to Alexandrov spaces and to their open subsets (in cases both of curvature bounded below and curvature bounded above). As we shall see, in all these cases LinDim behaves nicely and is easy to work with.

Recall that a cone map is a map between cones respecting the cone multiplication.

7.9. Definition of linear dimension. *Let $\mathcal{X}$ be a metric space with defined angles. The linear dimension of $\mathcal{X}$ (denoted by* $\operatorname{LinDim}\mathcal{X}$*) is defined as the exact upper bound on $m \in \mathbb{Z}_{\geqslant 0}$ such that there is a distance-preserving cone embedding $\mathbb{E}^m \hookrightarrow \mathrm{T}_p\mathcal{X}$ for some $p \in \mathcal{X}$; here $\mathbb{E}^m$ denotes the m-dimensional Euclidean space and $\mathrm{T}_p\mathcal{X}$ denotes the tangent space of $\mathcal{X}$ at p (defined in Section 6D).*

Note that LinDim takes values in $\mathbb{Z}_{\geqslant 0} \cup \{\infty\}$.

The linear dimension LinDim has no immediate relations to HausDim and TopDim. Also, LinDim does not satisfy the cover axiom (7.4). Note that

❶ $$\operatorname{LinDim}(\mathcal{X} \times \mathcal{Y}) = \operatorname{LinDim}\mathcal{X} + \operatorname{LinDim}\mathcal{Y}$$

for any two metric spaces $\mathcal{X}$ and $\mathcal{Y}$ with defined angles.

The following exercise is based on a construction of Thomas Foertsch and Viktor Schroeder [120]; it shows that the condition on angles in ❶ cannot be removed.

7.10. Exercise. *Construct metrics* ρ_1 *and* ρ_2 *on* $\mathbb{R}^{10}$ *defined by norms, such that* $(\mathbb{R}^{10}, \rho_i)$ *do not contain an isometric copy of* $\mathbb{E}^2$ *but* $(\mathbb{R}^{10}, \rho_1) \times \times (\mathbb{R}^{10}, \rho_2)$ *has an isometric copy of* $\mathbb{E}^{10}$.

Remarks. Linear dimension was first introduced by Conrad Plaut in [112] under the name local dimension. Geometric dimension, introduced by Bruce Kleiner [77] is closely related; it coincides with the linear dimension for CBB and CAT spaces.

One can extend the definition to arbitrary metric spaces. To do this one should modify the definition of tangent space and take an arbitrary n-dimensional Banach space instead of the Euclidean n-space. For Alexandrov spaces (either CBB or CAT) this modification is equivalent to the definition we use in this book.

Part II

Fundamentals

Chapter 8

Fundamentals of curvature bounded below

A Four-point comparison

Recall (Section 6B) that the model angle $\tilde\measuredangle^\kappa(p\,{}^x_y)$ is defined if

$$|p-x|+|p-y|+|x-y|<\varpi^\kappa.$$

8.1. Four-point comparison. *A quadruple of points p,x^1,x^2,x^3 in a metric space satisfies* CBB(κ) *comparison if*

❶ $$\tilde\measuredangle^\kappa(p\,{}^{x^1}_{x^2})+\tilde\measuredangle^\kappa(p\,{}^{x^2}_{x^3})+\tilde\measuredangle^\kappa(p\,{}^{x^3}_{x^1})\leqslant 2\cdot\pi.$$

or at least one of the model angles $\tilde\measuredangle^\kappa(p\,{}^{x^i}_{x^j})$ is not defined.

8.2. Definition. *Let $\mathcal{L}$ be a metric space.*

a) $\mathcal{L}$ is CBB(κ) *if any quadruple $p,x^1,x^2,x^3\in\mathcal{L}$ satisfies* CBB(κ) *comparison.*

b) $\mathcal{L}$ is locally CBB(κ) *if any point $q\in\mathcal{L}$ admits a neighborhood $\Omega\ni q$ such that any quadruple $p,x^1,x^2,x^3\in\Omega$ satisfies* CBB(κ) *comparison.*

c) $\mathcal{L}$ is a CBB *space if $\mathcal{L}$ is* CBB(κ) *for some $\kappa\in\mathbb{R}$.*

Remarks

- ⋄ CBB(κ) length spaces are often called spaces with curvature $\geqslant\kappa$ in the sense of Alexandrov. These spaces will usually be denoted by $\mathcal{L}$, for $\mathcal{L}$ower curvature bound.

- ◇ In the definition of CBB(κ), when $\kappa > 0$ most authors assume in addition that the diameter is at most ϖ^κ. For a complete length space, the latter means that it is not isometric to one of the exceptional spaces, see 8.43. We do not make this assumption. In particular, we consider the real line to have curvature $\geqslant 1$.
- ◇ If $\kappa < K$, then any complete length CBB(K) space is CBB(κ). Moreover directly from the definition it follows that if $K \leqslant \leqslant 0$, then any CBB(K) space is CBB(κ). However, in the case $K > 0$ the latter statement does not hold and the former statement is not trivial; it will be proved in 8.32.

8.3. Exercise. *Let $\mathcal{L}$ be a metric space and $\kappa \leqslant 0$. Show that $\mathcal{L}$ is* CBB(κ) *if for any quadruple of points $p, x^1, x^2, x^3 \in \mathcal{L}$ there is a quadruple of points $q, y^1, y^2, y^3 \in \mathbb{M}^2(\kappa)$ such that*

$$|p - x^i| = |q - y^i| \quad \textit{and} \quad |x^i - x^j| \leqslant |y^i - y^j|$$

for all i and j.

The exercise above is a special case of $(1+n)$-point comparison (10.8).

8.4. Exercise. *Let $\mathcal{L}$ be a metric space. Show that $\mathcal{L}$ is* CBB(0) *if and only if*

$$\operatorname{area}\tilde{\triangle}^0(xyz) \leqslant \operatorname{area}\tilde{\triangle}^0(pxy) + \operatorname{area}\tilde{\triangle}^0(pyz) + \operatorname{area}\tilde{\triangle}^0(pzx)$$

for any 4 distinct points $p, x, y, z \in \mathcal{L}$.

Recall that ω denotes a selective ultrafilter on $\mathbb{N}$, which is fixed once and for all. The following proposition follows directly from the definition of CBB(κ) comparison and the definitions of ω-limit and ω-power given in Section 3B.

8.5. Proposition. *Let $\mathcal{L}_n$ be a* CBB(κ_n) *space for each n. Assume $\mathcal{L}_n \to \mathcal{L}_\omega$ and $\kappa_n \to \kappa_\omega$ as $n \to \omega$. Then $\mathcal{L}_\omega$ is* CBB(κ_ω).

Moreover, a metric space $\mathcal{L}$ is CBB(κ) *if and only if so is its ultrapower $\mathcal{L}^\omega$.*

8.6. Theorem. *Let $\mathcal{L}$ be a* CBB(κ) *space, $\mathcal{M}$ be a metric space, and $\sigma\colon \mathcal{L} \to \mathcal{M}$ be a submetry. Assume p, x^1, x^2, x^3 is a quadruple of points in $\mathcal{M}$ such that $|p - x^i| < \frac{\varpi^\kappa}{2}$ for any i. Then the quadruple satisfies* CBB(κ) *comparison.*

In particular,

a) The space $\mathcal{M}$ is locally CBB(κ). *Moreover, any open ball of radius $\frac{\varpi^\kappa}{4}$ in $\mathcal{M}$ is* CBB(κ).

b) If $\kappa \leqslant 0$, then $\mathcal{M}$ is CBB(κ).

Corollary 8.33 gives a stronger statement; it states that if $\mathcal{L}$ is a complete length space, then $\mathcal{M}$ is always CBB(κ). The theorem above together with Proposition 5.8 imply the following:

8.7. Corollary. *Assume that the group G acts isometrically on a* CBB(κ) *space $\mathcal{L}$ and has closed orbits. Then the quotient space $\mathcal{L}/G$ is locally* CBB(κ).

Proof of 8.6. Fix a quadruple of points $p, x^1, x^2, x^3 \in \mathcal{M}$ such that $|p - x^i|_< \frac{\varpi^\kappa}{2}$ for any i. Choose an arbitrary $\hat{p} \in \mathcal{L}$ such that $\sigma(\hat{p}) = p$.

Since σ is submetry, we can choose the points $\hat{x}^1, \hat{x}^2, \hat{x}^3 \in \mathcal{L}$ such that $\sigma(\hat{x}_i) = x_i$ and

$$|p - x^i|_{\mathcal{M}} \lesseqgtr |\hat{p} - \hat{x}^i|_{\mathcal{L}} \pm \delta$$

for all i and any fixed $\delta > 0$.

Note that

$$|x^i - x^j|_{\mathcal{M}} \leqslant |\hat{x}^i - \hat{x}^j|_{\mathcal{L}} \leqslant |p - x^i|_{\mathcal{M}} + |p - x^j|_{\mathcal{M}} + 2{\cdot}\delta$$

for all i and j.

Since $|p - x^i| < \frac{\varpi^\kappa}{2}$, we can choose $\delta > 0$ above so that the angles $\tilde{\measuredangle}^\kappa(\hat{p}\,{}^{\hat{x}^i}_{\hat{x}^j})$ are defined. Moreover, given $\varepsilon > 0$, the value δ can be chosen in such a way that the inequality

❷ $$\tilde{\measuredangle}^\kappa(p\,{}^{x^i}_{x^j}) < \tilde{\measuredangle}^\kappa(\hat{p}\,{}^{\hat{x}^i}_{\hat{x}^j}) + \varepsilon$$

holds for all i and j.

By CBB(κ) comparison in $\mathcal{L}$, we have

$$\tilde{\measuredangle}^\kappa(\hat{p}\,{}^{\hat{x}^1}_{\hat{x}^2}) + \tilde{\measuredangle}^\kappa(\hat{p}\,{}^{\hat{x}^2}_{\hat{x}^3}) + \tilde{\measuredangle}^\kappa(\hat{p}\,{}^{\hat{x}^3}_{\hat{x}^1}) \leqslant 2{\cdot}\pi.$$

Applying ❷, we get

$$\tilde{\measuredangle}^\kappa(p\,{}^{x^1}_{x^2}) + \tilde{\measuredangle}^\kappa(p\,{}^{x^2}_{x^3}) + \tilde{\measuredangle}^\kappa(p\,{}^{x^3}_{x^1}) < 2{\cdot}\pi + 3{\cdot}\varepsilon.$$

Since $\varepsilon > 0$ is arbitrary we have

$$\tilde{\measuredangle}^\kappa(p\,{}^{x^1}_{x^2}) + \tilde{\measuredangle}^\kappa(p\,{}^{x^2}_{x^3}) + \tilde{\measuredangle}^\kappa(p\,{}^{x^3}_{x^1}) \leqslant 2{\cdot}\pi;$$

that is, the CBB(κ) comparison holds for this quadruple. □

B Geodesics

We are going to show that all complete length CBB spaces have plenty of geodesics in the following sense. Recall that a subset of a topological space is called G-delta if it is a countable intersection of open sets.

8.8. Definition. *A metric space $\mathcal{X}$ is called G-delta geodesic if for any point $p \in \mathcal{X}$ there is a dense G-delta set $W_p \subset \mathcal{X}$ such that for any $q \in W_p$ there is a geodesic $[pq]$.*

A metric space $\mathcal{X}$ is called locally G-delta geodesic if for any point $p \in \mathcal{X}$ there is a G-delta set $W_p \subset \mathcal{X}$ such that W_p is dense in a neighborhood of p and for any $q \in W_p$ there is a geodesic $[pq]$.

Recall that general complete length spaces might have no geodesics; see Exercise 2.10.

8.9. Exercise. *Construct a complete length* CBB(0) *space that is not geodesic.*

8.10. Definition. *Let $\mathcal{X}$ be a metric space and $p \in \mathcal{X}$. A point $q \in \mathcal{X}$ is called p-straight (briefly, $q \in \operatorname{Str}(p)$) if*

$$\varlimsup_{r\to q} \frac{|p-r| - |p-q|}{|q-r|} = 1.$$

For an array of points $x^1, x^2, \dots, x^k$, we use the notation

$$\operatorname{Str}(x^1, x^2, \dots, x^k) = \bigcap_{i=1}^{k} \operatorname{Str}(x^i).$$

8.11. Theorem. *Let $\mathcal{L}$ be a complete length* CBB *space and $p \in \mathcal{L}$. Then the set $\operatorname{Str}(p)$ is a dense G-delta set, and for any $q \in \operatorname{Str}(p)$ there is a unique geodesic $[pq]$.*

In particular, $\mathcal{L}$ is G-delta geodesic.

The proof below is very close to the original proof given by Conrad Plaut [112, Th. 27].

Proof. Given a positive integer n, consider the set Ω_n of all points $q \in \mathcal{L}$ such that

$$(1 - \tfrac{1}{n})\cdot|q-r| < |p-r| - |p-q| < \tfrac{1}{n}$$

for some $r \in \mathcal{L}$. Clearly Ω_n is open; let us show that Ω_n is dense in $\mathcal{L}$.

Assuming the contrary, there is a point $x \in \mathcal{L}$ such that

$$\mathrm{B}(x, \varepsilon) \cap \Omega_n = \varnothing$$

for some $\varepsilon > 0$. Since $\mathcal{L}$ is a length space, for any $\delta > 0$, there exists a point $y \in \mathcal{L}$ such that

$$|x-y| < \tfrac{\varepsilon}{2} + \delta \quad \text{and} \quad |p-y| < |p-x| - \tfrac{\varepsilon}{2} + \delta.$$

If ε and δ are sufficiently small, then

$$(1-\tfrac{1}{n})\cdot|y-x| < |p-x| - |p-y| < \tfrac{1}{n};$$

that is, $y \in \Omega_n$, a contradiction.

Note that $\mathrm{Str}(p) = \bigcap_{n\in\mathbb{N}} \Omega_n$. It follows that $\mathrm{Str}(p)$ is a dense G-delta set.

Assuming $q \in \mathrm{Str}(p)$, let us show that there is a unique geodesic connecting p and q. Note that it is sufficient to show that for all sufficiently small $t > 0$ there is a unique point z such that

❶ $$t = |q-z| = |p-q| - |p-z|.$$

First let us show uniqueness. Assume z and z' both satisfy ❶. Take a sequence $r_n \to q$ such that

$$\frac{|p-r_n| - |p-q|}{|q-r_n|} \to 1.$$

By the triangle inequality,

$$|z-r| - |z-q|, \quad |z'-r| - |z'-q| \geqslant |p-r| - |p-q|;$$

thus, as $n \to \infty$,

$$\frac{|z-r_n| - |z-q|}{|q-r_n|}, \quad \frac{|z'-r_n| - |z'-q|}{|q-r_n|} \to 1.$$

Therefore $\measuredangle^{\kappa}(q\,{}^{z}_{r_n}) \to \pi$ and $\measuredangle^{\kappa}(q\,{}^{z'}_{r_n}) \to \pi$. (Here we use that t is small, otherwise if $\kappa > 0$ the angles might be undefined.)

From $\mathrm{CBB}(\kappa)$ comparison (8.2), $\measuredangle^{\kappa}(q\,{}^{z}_{z'}) = 0$ and thus $z = z'$.

The proof of existence is similar. Choose a sequence r_n as above. Since $\mathcal{L}$ is a complete length space, there is a sequence $z_{\hbar} \in \mathcal{L}$ such that $|q-z_{\hbar}| \to t$ and $|p-q| - |p-z_{\hbar}| \to t$ as $\hbar \to \infty$. Then

$$\lim_{n\to\infty} \lim_{\hbar\to\infty} \measuredangle^{\kappa}(q\,{}^{z_{\hbar}}_{r_n}) = \pi.$$

Thus, for any $\varepsilon > 0$ and sufficiently large $n, \hbar$, we have $\measuredangle^{\kappa}(q\,{}^{z_{\hbar}}_{r_n}) > \pi - \varepsilon$. From $\mathrm{CBB}(\kappa)$ comparison (8.2), for all large $\hbar$ and j, we have $\measuredangle^{\kappa}(q\,{}^{z_{\hbar}}_{z_j}) < 2\cdot\varepsilon$ and thus

$$|z_{\hbar} - z_j| < \varepsilon\cdot c(\kappa, t);$$

that is, $\{z_n\}$ is Cauchy, and $z = \lim_n z_n$ satisfies ❶. □

8.12. Exercise. *Let $\mathcal{L}$ be a complete length* CBB *space and $A \subset \mathcal{L}$ be a closed subset. Show that there is a dense G-delta set $W \subset \mathcal{L}$ such that* *for any $q \in W$, there is a unique geodesic $[pq]$ with $p \in A$ that realizes the distance from q to A; that is, $|p - q| = \operatorname{dist}_A q$.*

8.13. Exercise. *Construct a complete length* CBB *space $\mathcal{L}$ with an everywhere dense G-delta set A such that $A \cap]xy[= \varnothing$ for any geodesic $[xy]$ in $\mathcal{L}$.*

C More comparisons

The following theorem makes it easier to use Euclidean intuition in the Alexandrov setting.

8.14. Theorem. *If $\mathcal{L}$ is a* CBB(κ) *space, then the following conditions hold for all $p, x, y \in \mathcal{L}$, provided the model triangle $\tilde{\triangle}^\kappa(pxy)$ is defined.*

a) (adjacent angle comparison) for any geodesic $[xy]$ and $z \in]xy[$, $z \neq p$ we have

$$\tilde{\measuredangle}^\kappa(z\,{}^p_x) + \tilde{\measuredangle}^\kappa(z\,{}^p_y) \leqslant \pi.$$

b) (point-on-side comparison) for any geodesic $[xy]$ and $z \in]xy[$, we have

$$\tilde{\measuredangle}^\kappa(x\,{}^p_y) \leqslant \tilde{\measuredangle}^\kappa(x\,{}^p_z);$$

or, equivalently,

$$|\tilde{p} - \tilde{z}| \leqslant |p - z|,$$

where $[\tilde{p}\tilde{x}\tilde{y}] = \tilde{\triangle}^\kappa(pxy)$, $\tilde{z} \in]\tilde{x}\tilde{y}[$, $|\tilde{x} - \tilde{z}| = |x - z|$.

c) (hinge comparison) for any hinge $[x\,{}^p_y]$, the angle $\measuredangle[x\,{}^p_y]$ is defined and

$$\measuredangle[x\,{}^p_y] \geqslant \tilde{\measuredangle}^\kappa(x\,{}^p_y),$$

or equivalently

$$\tilde{\curlyvee}^\kappa[x\,{}^p_y] \geqslant |p - y|.$$

Moreover, if $z \in]xy[$ and $z \neq p$, then

$$\measuredangle[z\,{}^p_y] + \measuredangle[z\,{}^p_x] \leqslant \pi$$

for any two hinges $[z\,{}^p_y]$ and $[z\,{}^p_x]$ with common side $[zp]$.

Moreover, in each case, the converse holds if $\mathcal{L}$ is G-delta geodesic. That is, if one of the conditions (a), (b) or (c) holds in a G-delta geodesic space $\mathcal{L}$, then $\mathcal{L}$ is CBB(κ)*.*

Remarks. Monotonicity of the model angle with respect to adjacent sidelengths (8.17) was named the c o n v e x i t y p r o p e r t y by Alexandrov.

A slightly stronger form of (c) is given in 8.28. See also Problem 8.50.

Proof; (a). Since $z\in\,]xy[$, we have $\tilde\measuredangle^\kappa(z\,{}^x_y)=\pi$. Thus, CBB($\kappa$) comparison

$$\tilde\measuredangle^\kappa(z\,{}^x_y)+\tilde\measuredangle^\kappa(z\,{}^p_x)+\tilde\measuredangle^\kappa(z\,{}^p_y)\leqslant 2\cdot\pi$$

implies

$$\tilde\measuredangle^\kappa(z\,{}^p_x)+\tilde\measuredangle^\kappa(z\,{}^p_y)\leqslant\pi.$$

(a) ⇔ *(b).* Follows from Alexandrov's lemma (6.2).

(a) + (b) ⇒ *(c).* From (b) we get that for $\bar p\in\,]xp]$ and $\bar y\in\,]xy]$, the function $(|x-\bar p|,|x-\bar y|)\mapsto\tilde\measuredangle^\kappa(x\,{}^{\bar p}_{\bar y})$ is nonincreasing in each argument. In particular, $\measuredangle[x\,{}^p_y]=\sup\{\tilde\measuredangle^\kappa(x\,{}^{\bar p}_{\bar y})\}$ is defined and is at least $\tilde\measuredangle^\kappa(x\,{}^p_y)$.

From above and (a), it follows that

$$\measuredangle[z\,{}^p_y]+\measuredangle[z\,{}^p_x]\leqslant\pi.$$

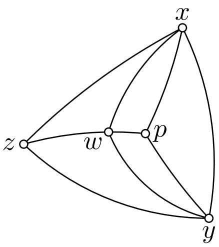

Converse. Assume first that $\mathcal{L}$ is geodesic. Consider a point $w\in\,]pz[$ close to p. From (c), it follows that

$$\measuredangle[w\,{}^x_z]+\measuredangle[w\,{}^x_p]\leqslant\pi\quad\text{and}\quad\measuredangle[w\,{}^y_z]+\measuredangle[w\,{}^y_p]\leqslant\pi.$$

Since $\measuredangle[w\,{}^x_y]\leqslant\measuredangle[w\,{}^x_p]+\measuredangle[w\,{}^y_p]$ (see 6.4), we get

$$\measuredangle[w\,{}^x_z]+\measuredangle[w\,{}^y_z]+\measuredangle[w\,{}^x_y]\leqslant 2\cdot\pi.$$

Applying the first inequality in (c),

$$\tilde\measuredangle^\kappa(w\,{}^x_z)+\tilde\measuredangle^\kappa(w\,{}^y_z)+\tilde\measuredangle^\kappa(w\,{}^x_y)\leqslant 2\cdot\pi.$$

Passing to the limits $w\to p$, we have

$$\tilde\measuredangle^\kappa(p\,{}^x_z)+\tilde\measuredangle^\kappa(p\,{}^y_z)+\tilde\measuredangle^\kappa(p\,{}^x_y)\leqslant 2\cdot\pi.$$

If $\mathcal{L}$ is only G-delta geodesic, we can apply the above arguments to sequences of points $p_n,w_n\to p$, $x_n\to x$, $y_n\to y$ such that $[p_nz]$ exists, $w_n\in\,]zp_n[$ and $[x_nw_n]$, $[y_nw_n]$ exist, and then pass to the limit as $n\to\to\infty$. □

8.15. Exercise. *Let $\mathcal{L}$ be $\mathbb{R}^m$ with a metric defined by a norm. Show that $\mathcal{L}$ is a complete length* CBB *space if and only if $\mathcal{L}\overset{iso}{=\!=}\mathbb{E}^m$.*

8.16. Exercise. *Assume $\mathcal{L}$ is a complete length* CBB *space, and $[px]$, $[py]$ be two geodesics in the same geodesic direction $\xi\in\Sigma'_p$. Show that*

$$[px]\subset[py]\quad\textit{or}\quad[px]\supset[py].$$

8.17. Angle-sidelength monotonicity. *Let $\mathcal{L}$ be a complete length* CBB(κ) *space, $p, x, y \in \mathcal{L}$, $\tilde\triangle^\kappa(pxy)$ is defined and there is a geodesic $[xy]$. Then for $\bar y \in \,]xy]$ the function*

$$|x - \bar y| \mapsto \tilde\measuredangle^\kappa(x\,{}^p_{\bar y})$$

is nonincreasing.

In particular, if a geodesic $[xp]$ exists and $\bar p \in \,]xp]$, then

a) the function

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

is nonincresing in each argument

b) The angle $\measuredangle[x\,{}^p_y]$ is defined and

$$\measuredangle[x\,{}^p_y] = \sup\left\{\, \tilde\measuredangle^\kappa(x\,{}^{\bar p}_{\bar y}) \,:\, \bar p \in \,]xp],\ \bar y \in \,]xy] \,\right\}.$$

The proof is contained in the first part of (a)+(b)⇒(c) of Theorem 8.14.

8.18. Exercise. *Let $\mathcal{L}$ be a complete length* CBB(κ) *space, $p, x, y \in \mathcal{L}$ and $v, w \in \,]xy[$. Prove that*

$$\tilde\measuredangle^\kappa(x\,{}^y_p) = \tilde\measuredangle^\kappa(x\,{}^v_p) \quad\iff\quad \tilde\measuredangle^\kappa(x\,{}^y_p) = \tilde\measuredangle^\kappa(x\,{}^w_p).$$

8.19. Advanced exercise. *Construct a geodesic space $\mathcal{X}$ that is not* CBB(0)*, but meets the following condition: for any 3 points $p, x, y \in \mathcal{X}$ there is a geodesic $[xy]$ such that for any $z \in \,]xy[$*

$$\tilde\measuredangle^0(z\,{}^p_x) + \tilde\measuredangle^0(z\,{}^p_y) \leqslant \pi.$$

8.20. Advanced exercise. *Let $\mathcal{L}$ be a complete length space such that for any quadruple $p, x, y, z \in \mathcal{L}$ the following inequality holds*

❶ $$|p - x|^2 + |p - y|^2 + |p - z|^2 \geqslant \tfrac{1}{3}\cdot\left[|x - y|^2 + |y - z|^2 + |z - x|^2\right].$$

Prove that $\mathcal{L}$ is CBB(0)*.*

Construct a 4-point metric space $\mathcal{X}$ that satisfies inequality ❶ *for any relabeling of its points by p, x, y, z, such that $\mathcal{X}$ is not* CBB(0)*.*

Assume that for a given triangle $[x^1x^2x^3]$ in a metric space its κ-model triangle $[\tilde x^1\tilde x^2\tilde x^3] = \tilde\triangle^\kappa(x^1x^2x^3)$ is defined. We say the triangle $[x^1x^2x^3]$ is κ-thick if the natural map (see definition 9.19) $[\tilde x^1\tilde x^2\tilde x^3] \to [x^1x^2x^3]$ is distance non contracting.

8.21. Exercise. *Prove that any triangle with perimeter $< \varpi^\kappa$ in a* CBB(κ) *space is* κ*-thick.*

The following exercise is inspired by Busemann's definition [35].

8.22. Exercise.

a) Show that any CBB(0) *space* $\mathcal{L}$ *satisfies the following condition: for any three points* $p, q, r \in \mathcal{L}$*, if* $\bar{q}$ *and* $\bar{r}$ *are midpoints of geodesics* $[pq]$ *and* $[pr]$ *respectively, then* $2{\cdot}|\bar{q}-\bar{r}| \geqslant |q-r|$*.*

b) Show that there is a metric on $\mathbb{R}^2$ *defined by a norm that satisfies the above condition, but is not* CBB(0)*.*

D Function comparison

In this section we will translate the angle comparison definitions (Theorem 8.14) to a concavity-like property of the distance functions as defined in Section 5F. This is a conceptual step — we reformulate a global geometric condition into an infinitesimal condition on distance functions.

8.23. Theorem. *Let* $\mathcal{L}$ *be a complete length space. Then the following statements are equivalent:*

a) $\mathcal{L}$ *is* CBB(κ)*.*

b) (function comparison) $\mathcal{L}$ *is G-delta geodesic and for any* $p \in \mathcal{L}$*, the function* $f = \operatorname{md}^\kappa \circ \operatorname{dist}_p$ *satisfies the differential inequality*

$$f'' \leqslant 1 - \kappa{\cdot}f.$$

in $\mathrm{B}(p, \varpi^\kappa)$*.*

In particular, a complete G-delta geodesic space $\mathcal{L}$ is CBB(0) if and only if for any $p \in \mathcal{L}$, the function $\operatorname{dist}_p^2 : \mathcal{L} \to \mathbb{R}$ is 2-concave as defined in Section 5F.

Proof. Let $[xy]$ be a geodesic in $\mathrm{B}(p, \varpi^\kappa)$ and $\ell = |x-y|$. Consider the model triangle $[\tilde{p}\tilde{x}\tilde{y}] = \tilde{\triangle}^\kappa(pxy)$. Set

$$\tilde{r}(t) = |\tilde{p} - \operatorname{geod}_{[\tilde{x}\tilde{y}]}(t)|, \qquad r(t) = |p - \operatorname{geod}_{[xy]}(t)|.$$

Clearly $\tilde{r}(0) = r(0)$ and $\tilde{r}(\ell) = r(\ell)$. Set $\tilde{f} = \operatorname{md}^\kappa \circ \tilde{r}$ and $f = \operatorname{md}^\kappa \circ r$. From 1.1*a* we get that $\tilde{f}'' = 1 - \kappa{\cdot}\tilde{f}$.

Note that the point-on-side comparison (8.14*b*) for point p and geodesic $[xy]$ is equivalent to $\tilde{r} \leqslant r$. Since md^κ is increasing on $[0, \varpi^\kappa)$, $\tilde{r} \leqslant r$ is equivalent to $\tilde{f} \leqslant f$. The latter is Jensen's inequality (5.14*c*) for the function $t \mapsto \operatorname{md}^\kappa|p - \operatorname{geod}_{[xy]}(t)|$ on the interval $[0, \ell]$. Hence the result. □

Recall that Busemann functions are defined in Proposition 6.1. Compare the following exercise to Exercise 9.27.

8.24. Exercise. *Let* $\mathcal{L}$ *be a complete length* CBB(κ) *space and* $\operatorname{bus}_\gamma : \mathcal{L} \to \mathbb{R}$ *be the Busemann function for a half-line* $\gamma : [0, \infty) \to \mathcal{L}$*.*

a) If $\kappa = 0$*, then the Busemann function* $\operatorname{bus}_\gamma$ *is concave.*

b) If $\kappa = -1$*, then the function*

$$f = \exp \circ \operatorname{bus}_\gamma$$

satisfies

$$f'' - f \leqslant 0.$$

E Development

In this section we reformulate the function comparison using a more geometric language based on the definition of development given below.

This definition appears in [11] and an earlier form of it can be found in [86]. The definition is somewhat lengthy, but it defines a useful comparison object for a curve. Often it is easier to write proofs in terms of function comparison but think in terms of developments.

8.25. Lemma-definition. *Let* $\kappa \in \mathbb{R}$*,* $\mathcal{X}$ *be a metric space,* $\gamma\colon \mathbb{I} \to \mathcal{X}$ *be a* 1*-Lipschitz curve,* $p \in \mathcal{X}$*,* $\tilde p \in \mathbb{M}^2(\kappa)$*. Assume* $0 < |p - \gamma(t)| < \varpi^\kappa$ *for all* $t \in \mathbb{I}$*. Then there exists a unique up to rotation curve* $\tilde\gamma\colon \mathbb{I} \to \mathbb{M}^2(\kappa)$*, 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 $p, \tilde p, \gamma, \tilde\gamma$ *are as above, then* $\tilde\gamma$ *is called the* κ*-development of* γ *with respect to* p*; the point* $\tilde p$ *is called the basepoint of the development. When we say that the* κ*-development of* γ *with respect to* p *is defined we always assume that* $0 < |p - \gamma(t)| < \varpi^\kappa$ *for all* $t \in \mathbb{I}$*.*

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

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

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 nondecreasing.

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

We need the following analogs of sub- and super-graphs and convex/concave functions, adapted to polar coordinates in $\mathbb{M}^2(\kappa)$.

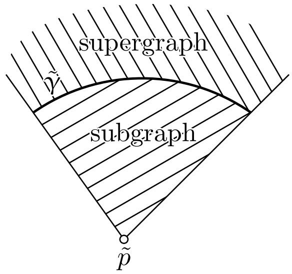

8.26. Definition. *Let $\tilde\gamma\colon \mathbb{I}\to\mathbb{M}^2(\kappa)$ be a curve and $\tilde p\in\mathbb{M}^2(\kappa)$ be such that there is a unique geodesic $[\tilde p\tilde\gamma(t)]$ for any $t\in\mathbb{I}$ and the direction of $[\tilde p\tilde\gamma(t)]$ turns monotonically as t grows.*

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 $\tilde x\in\mathbb{M}^2(\kappa)$ 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.

The curve $\tilde\gamma$ is called locally convex (concave) with respect to $\tilde p$ if for any interior value t_0 in $\mathbb{I}$ there is a subsegment $(a,b)\subset\mathbb{I}$, $(a,b)\ni$ $\ni t_0$, such that the restriction $\tilde\gamma|_{(a,b)}$ is convex (concave) with respect to $\tilde p$.

Note that if $\kappa>0$, then the supergraph of a curve is the subgraph with respect to the opposite point.

For developments, all the notions above will be considered with respect to their basepoints. In particular, if $\tilde\gamma$ is a development, we will say it is (locally) convex if it is (locally) convex with respect to its basepoint.

8.27. Development comparison. *A complete G-delta geodesic space $\mathcal{L}$ is* CBB(κ) *if and only if for any point $p\in\mathcal{L}$ and any geodesic γ in* $\mathrm{B}(p,\varpi^\kappa)\backslash\{p\}$*, its κ-development with respect to p is convex.*

A simpler proof of the "only-if"-part can be built on the adjacent angle comparison (8.14*a*). We use a longer proof since it also implies the short hinge lemma (8.28).

Proof; "only-if"-part. Let $\gamma\colon[0,T]\to\mathrm{B}(p,\varpi^\kappa)\backslash\{p\}$ be a unit-speed geodesic in $\mathcal{L}$.

Consider a fine partition

$$0=t_0<t_1<\ldots<t_n=T.$$

Set $x_i=\gamma(t_i)$ and choose a point

$$p'\in\mathrm{Str}(x_0,x_1,\ldots,x_n)$$

sufficiently close to p; so geodesics $[p'x_i]$ exist for all i (see Definition 8.10).

Let us construct a chain of model triangles $[\tilde p'\tilde x_{i-1}\tilde x_i]=\tilde\triangle^\kappa(p'x_{i-1}x_i)$ in such a way that direction $[\tilde p'\tilde x_i]$ turns counterclockwise as i grows. By the hinge comparison (8.14c), we have

❶
$$\begin{aligned}\measuredangle[\tilde x_i\,{}^{\tilde x_{i-1}}_{\tilde p'}]+\measuredangle[\tilde x_i\,{}^{\tilde x_{i+1}}_{\tilde p'}]&=\tilde\measuredangle^\kappa(x_i\,{}^{x_{i-1}}_{p'})+\tilde\measuredangle^\kappa(x_i\,{}^{x_{i+1}}_{p'})\leqslant\\&\leqslant\measuredangle[x_i\,{}^{x_{i-1}}_{p'}]+\measuredangle[x_i\,{}^{x_{i+1}}_{p'}]\leqslant\\&\leqslant\pi.\end{aligned}$$

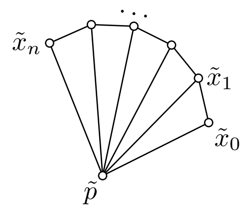

Further, since γ is a unit-speed geodesic, we have

❷
$$\sum_{i=1}^n|x_{i-1}-x_i|\leqslant|p'-x_0|+|p'-x_n|.$$

Since $p'\notin\gamma$, the development comparison implies that $\tilde p'$ does not lie on the polygonal line $\tilde x_0\dots\tilde x_n$.

If $\kappa\leqslant 0$, then ❷ implies that

❸
$$\vartheta:=\sum_{i=1}^n\measuredangle[\tilde p'\,{}^{\tilde x_i}_{\tilde x_{i-1}}]\leqslant\pi.$$

In the case $\kappa>0$, the proof of ❸ requires more work. Applying rescaling, we can assume that $\kappa=1$. Since γ lies in $B_\pi(p')$, the point-on-side comparison implies that the antipodal point of $\tilde p'$ does not lie on the polygonal line $\tilde x_0\dots\tilde x_n$.

Consider the space L glued from the model triangles $[\tilde p'\tilde x_0\tilde x_1],\dots$ $\dots,[\tilde p'\tilde x_{n-1}\tilde x_n]$ along the corresponding sides. Note that ϑ is the total angle of L at $\tilde p'$. We can assume that L has nonempty interior. Otherwise all the triangles are degenerate. Since $\tilde p'$ is not on the polygonal line $\tilde x_0\dots\tilde x_n$, the latter implies that $\vartheta=0$.

Consider a minimizing geodesic $[\tilde x_0\tilde x_n]_L$. By ❷ we may assume that $\tilde p'\notin[\tilde x_0\tilde x_n]_L$. Further if the geodesic $[\tilde x_0\tilde x_n]_L$ contains one of the points $\tilde x_1,\dots,\tilde x_{n-1}$, then it coincides with the polygonal line $\tilde x_0\dots\tilde x_n$. (In particular we have equality in ❶ for each i.) In this case, the sum in the left-hand side of ❷ must be at most π; otherwise $[\tilde x_0\tilde x_n]_L$ is not minimizing. Therefore ❸ follows. In the remaining case $[\tilde x_0\tilde x_n]_L$ meets the boundary of L only at its ends. In this case, $|\tilde x_0-\tilde x_n|_L\leqslant\pi$; otherwise $[\tilde x_0\tilde x_n]_L$ is not minimizing. Whence ❸ follows.

Inequalities ❶ and ❸ imply that the polygon $[\tilde p'\tilde x_0\tilde x_1\dots\tilde x_n]$ is convex.

Let us take finer and finer partitions and pass to the limit of the polygon $\tilde p'\tilde x_0\tilde x_1\dots\tilde x_n$ as $p'\to p$. We obtain a convex curvilinear triangle formed by a curve $\tilde\gamma\colon[0,T]\to\mathbb M^2(\kappa)$ — the limit of broken line $\tilde x_0\tilde x_1\dots\tilde x_n$ and two geodesics $[\tilde p'\,\tilde\gamma(0)]$, $[\tilde p'\,\tilde\gamma(T)]$. Since $[\tilde p'\tilde x_0\tilde x_1\dots\tilde x_n]$ is convex, the natural parametrization of $\tilde x_0\tilde x_1\dots\tilde x_n$ converges to the natural parametrization of $\tilde\gamma$. Thus $\tilde\gamma$ is the κ-development of γ with respect to p. This proves the "only-if" part of (8.27).

Proof; "if"-part. Assuming convexity of the development, we will prove the point-on-side comparison (8.14b). We can assume that $p\notin[xy]$; otherwise the statement is trivial.

Set $T=|x-y|$ and $\gamma(t)=\mathrm{geod}_{[xy]}(t)$; note that γ is a geodesic in $\mathrm{B}(p,\varpi^\kappa)\backslash\{p\}$. Let $\tilde\gamma\colon[0,T]\to\mathbb{M}^2(\kappa)$ be the κ-development with base $\tilde p$ of γ with respect to p. Take a partition $0=t_0<t_1<\ldots<t_n=T$, and set

$$\tilde y_i=\tilde\gamma(t_i)\quad\text{and}\quad\tau_i=|\tilde y_0-\tilde y_1|+|\tilde y_1-\tilde y_2|+\ldots+|\tilde y_{i-1}-\tilde y_i|.$$

Since $\tilde\gamma$ is convex, for a fine partition we have that broken line $\tilde y_0\tilde y_1\ldots\tilde y_n$ is also convex. Applying Alexandrov's lemma (6.2) inductively to pairs of model triangles

$$\tilde\triangle^\kappa\{\tau_{i-1},|\tilde p-\tilde y_0|,|\tilde p-\tilde y_{i-1}|\}$$

and

$$\tilde\triangle^\kappa\{|\tilde y_{i-1}-\tilde y_i|,|\tilde p-\tilde y_{i-1}|,|\tilde p-\tilde y_i|\}$$

we obtain that the sequence $\tilde\measuredangle^\kappa\{|\tilde p-\tilde y_i|;|\tilde p-\tilde y_0|,\tau_i\}$ is non increasing.

For finer and finer partitions we have

$$\max_i\{|\tau_i-t_i|\}\to 0.$$

Thus, the point-on-side comparison (8.14b) follows. □

Note that in the proof of "if"-part we could use a slightly weaker version of the hinge comparison (8.14c). Namely we proved the following lemma, which will be needed later in the proof of the globalization theorem (8.30).

8.28. Short hinge lemma. *Let $\mathcal{L}$ be a complete G-delta geodesic space such that for any hinge $[x\,{}^p_y]$ in $\mathcal{L}$ the angle $\measuredangle[x\,{}^p_y]$ is defined, and moreover if $x\in\,]yz[$ then*

$$\measuredangle[x\,{}^p_y]+\measuredangle[x\,{}^p_z]\leqslant\pi.$$

Assume that for any hinge $[x\,{}^p_y]$ in $\mathcal{L}$ we have

$$|p-x|+|x-y|<\varpi^\kappa\quad\Rightarrow\quad\measuredangle[x\,{}^p_y]\geqslant\tilde\measuredangle^\kappa(x\,{}^p_y).$$

Then $\mathcal{L}$ is CBB(κ).

F Local definitions and globalization

In this section we discuss locally CBB(κ) spaces. In particular, we prove the globalization theorem: equivalence of local and global definitions for complete length spaces.

The following theorem summarizes equivalent definitions of locally CBB(κ) spaces

8.29. Theorem. *Let $\mathcal{X}$ be a complete length space and $p\in\mathcal{X}$. Then the following conditions are equivalent:*

1) *(local CBB(κ) comparison) there is $R_1>0$ such that the comparison*
$$\tilde\measuredangle^\kappa(q\,{}^{x^1}_{x^2})+\tilde\measuredangle^\kappa(q\,{}^{x^2}_{x^3})+\tilde\measuredangle^\kappa(q\,{}^{x^3}_{x^1})\leqslant 2\cdot\pi$$
holds for any $q,x^1,x^2,x^3\in \mathrm{B}(p,R_1)$.
2) *(local Kirszbraun property) there is $R_2>0$ such that for any 3-point subset F_3 and any 4-point subset $F_4\supset F_3$ in $\mathrm{B}(p,R_2)$, any short map $f\colon F_3\to\mathbb{M}^2(\kappa)$ can be extended to a short map $\bar f\colon F_4\to\mathbb{M}^2(\kappa)$ (so $f=\bar f|_{F_3}$).*
3) *(local function comparison) there is $R_3>0$ such that $\mathrm{B}(p,R_3)$ is G-delta geodesic and for any $q\in\mathrm{B}(p,R_3)$, the function $f=\mathrm{md}^\kappa\circ\mathrm{dist}_q$ satisfies $f''\leqslant 1-\kappa\cdot f$ in $\mathrm{B}(p,R_3)$.*
4) *(local adjacent angle comparison) there is $R_4>0$ such that $\mathrm{B}(p,R_4)$ is G-delta geodesic, and if q and a geodesic $[xy]$ lie in $\mathrm{B}(p,R_4)$ and $z\in\,]xy[$, then*
$$\tilde\measuredangle^\kappa(z\,{}^q_x)+\tilde\measuredangle^\kappa(z\,{}^q_y)\leqslant\pi.$$
5) *(local point-on-side comparison) there is $R_5>0$ such that $\mathrm{B}(p,R_5)$ is G-delta geodesic and if q and a geodesic $[xy]$ lie in $\mathrm{B}(p,R_5)$ and $z\in\,]xy[$, then*
$$\tilde\measuredangle^\kappa(x\,{}^q_y)\leqslant\tilde\measuredangle^\kappa(x\,{}^q_z);$$
or, equivalently,
$$|\tilde p-\tilde z|\leqslant|p-z|,$$
where $[\tilde p\tilde x\tilde y]=\tilde\triangle^\kappa(pxy)$, $\tilde z\in\,]\tilde x\tilde y[$, $|\tilde x-\tilde z|=|x-z|$.
6) *(local hinge comparison) there is $R_6>0$ such that $\mathrm{B}(p,R_6)$ is G-delta geodesic and if $x\in\mathrm{B}(p,R_6)$, then*
(1) for any hinge $[x\,{}^q_y]$, the angle $\measuredangle[x\,{}^q_y]$ is defined, and
(2) if $x\in\,]yz[$, then[1]
$$\measuredangle[x\,{}^q_y]+\measuredangle[x\,{}^q_z]\leqslant\pi.$$
Moreover, if a hinge $[x\,{}^q_y]$ lies in $\mathrm{B}(p,R_6)$, then
$$\measuredangle[x\,{}^q_y]\geqslant\tilde\measuredangle^\kappa(x\,{}^q_y),$$
or, equivalently,
$$\tilde\curlyvee^\kappa[x\,{}^q_y]\geqslant|q-y|.$$

[1]Let us call that $[x\,{}^q_y]$ and $[x\,{}^q_z]$ are short notations for the pairs $([xq],[xy])$ and $([xq],[xz])$, thus these two hinges automatically have common side $[xq]$.

7) (local development comparison) there is $R_7 > 0$ such that $\mathrm{B}(p, R_7)$ is G-delta geodesic, and if a geodesic γ lies in $\mathrm{B}(p, R_7)$ and $q \in \mathrm{B}(p, R_7)\backslash\gamma$, then the κ-development $\tilde{\gamma}$ with respect to q is convex.

Moreover, for each pair $i, j \in \{1, 2, \dots, 7\}$ we can assume that

$$R_i > \tfrac{1}{9}\cdot R_j.$$

The proofs of each of these equivalences repeat the proofs of the corresponding global equivalences in localized form; see the proofs of Theorems 8.14, 8.23, 8.25, 10.1.

8.30. Globalization theorem. *Any complete length locally* $\mathrm{CBB}(\kappa)$ *space is* $\mathrm{CBB}(\kappa)$.

In the two-dimensional case this theorem was proved by Paolo Pizzetti [109]; later it was reproved independently by Alexandr Alexandrov [11]. Victor Toponogov [125] proved it for Riemannian manifolds of all dimensions. In the above generality, the theorem first appears in the paper of Michael Gromov, Yuriy Burago, and Grigory Perelman [34]; simplifications and modifications were given by Conrad Plaut [111], Katsuhiro Shiohama [123], and in the book of Dmitry Burago, Yuriy Burago, and Sergei Ivanov [27]. A generalization for non-complete but geodesic spaces was obtained by the third author [105]; namely it solves the following exercise:

8.31. Advanced exercise. *Assume $\mathcal{X}$ is a geodesic locally* $\mathrm{CBB}(\kappa)$ *space. Prove that the completion of $\mathcal{X}$ is* $\mathrm{CBB}(\kappa)$.

Our proof of Globalization theorem 8.30 is based on presentations in [111] and [27]; this proof was rediscovered independently by Urs Lang and Viktor Schroeder [80].

The following corollary of the globalization theorem says that the expression "space with curvature $\geqslant \kappa$" makes sense.

8.32. Corollary. *Let $\mathcal{L}$ be a complete length space. Then $\mathcal{L}$ is* $\mathrm{CBB}(\mathrm{K})$ *if and only if $\mathcal{L}$ is* $\mathrm{CBB}(\kappa)$ *for any* $\kappa < \mathrm{K}$.

Proof. Note that if $\mathrm{K} \leqslant 0$, this statement follows directly from the definition of an Alexandrov space (8.2) and monotonicity of the function $\kappa \mapsto \tilde{\measuredangle}^{\kappa}(x\,{}^{y}_{z})$ (1.1d).

The "if"-part also follows directly from the definition.

For $\mathrm{K} > 0$, the angle $\tilde{\measuredangle}^{\mathrm{K}}(x\,{}^{y}_{z})$ might be undefined while $\tilde{\measuredangle}^{\kappa}(x\,{}^{y}_{z})$ is defined. However, $\tilde{\measuredangle}^{\mathrm{K}}(x\,{}^{y}_{z})$ is defined if x, y, and z are sufficiently close to each other. Thus, if $\mathrm{K} > \kappa$, then any $\mathrm{CBB}(\mathrm{K})$ space is locally $\mathrm{CBB}(\kappa)$. It remains to apply the globalization theorem. $\square$

8.33. Corollary. *Let $\mathcal{L}$ be a complete length* CBB(κ) *space. Assume that a space $\mathcal{M}$ is the target space of a submetry from $\mathcal{L}$. Then $\mathcal{M}$ is a complete length space* CBB(κ) *space.*

In particular, if $G \curvearrowright \mathcal{L}$ is an isometric group action with closed orbits, then the quotient space $\mathcal{L}/G$ is a complete length CBB(κ) *space.*

Proof. This follows from the globalization theorem and Theorem 8.6. □

In the proof of the globalization theorem we will use three lemmas. The first is the short hinge lemma (8.28), which gives a characterization of CBB(κ) spaces that is the same as the hinge comparison condition 8.14*c* if $\kappa \leqslant 0$ and is a slightly weaker condition if $\kappa > 0$.

The following lemma says that if comparison holds for all small hinges, then it holds for slightly bigger hinges near the given point.

8.34. Key lemma. *Let $\kappa \in \mathbb{R}$, $0 < \ell \leqslant \varpi^\kappa$, $\mathcal{X}$ be a complete geodesic space and $p \in \mathcal{X}$ be a point such that $\mathrm{B}(p, 2\cdot\ell)$ is locally* CBB(κ).

Assume that for any point $q \in \mathrm{B}(p, \ell)$ the comparison

$$\measuredangle[x\,{}^{y}_{q}] \geqslant \tilde\measuredangle^\kappa(x\,{}^{y}_{q})$$

holds for any hinge $[x\,{}^{y}_{q}]$ with $|x-y|+|x-q| < \tfrac{2}{3}\cdot\ell$. Then the comparison

$$\measuredangle[x\,{}^{p}_{q}] \geqslant \tilde\measuredangle^\kappa(x\,{}^{p}_{q})$$

holds for any hinge $[x\,{}^{p}_{q}]$ with $|x-p|+|x-q| < \ell$.

Proof. It is sufficient to prove the inequality

❶ $$\tilde\Upsilon^\kappa[x\,{}^{p}_{q}] \geqslant |p-q|$$

for any hinge $[x\,{}^{p}_{q}]$ with $|x-p|+|x-q| < \ell$.

Fix q. Consider a hinge $[x\,{}^{p}_{q}]$ such that

$$\tfrac{2}{3}\cdot\ell \leqslant |p-x|+|x-q| < \ell.$$

First we construct a new smaller hinge $[x'\,{}^{p}_{q}]$ with

❷ $$|p-x|+|x-q| \geqslant |p-x'|+|x'-q|,$$

such that

❸ $$\tilde\Upsilon^\kappa[x\,{}^{p}_{q}] \geqslant \tilde\Upsilon^\kappa[x'\,{}^{p}_{q}].$$

Assume $|x-q| \geqslant |x-p|$; otherwise switch the roles of p and q in the following construction. Take $x' \in [xq]$ such that

❹ $$|p-x|+3\cdot|x-x'| = \tfrac{2}{3}\cdot\ell.$$

Choose a geodesic $[x'p]$ and consider the hinge $[x'\,{}^p_q]$ formed by $[x'p]$ and $[x'q] \subset [xq]$.[2] Then ❷ follows from the triangle inequality.

Further, note that we have $x, x' \in \mathrm{B}(p,\ell) \cap \mathrm{B}(q,\ell)$ and moreover

$$|p - x| + |x - x'| < \tfrac{2}{3}\cdot\ell, \qquad |p - x'| + |x' - x| < \tfrac{2}{3}\cdot\ell.$$

In particular,

❺ $$\measuredangle[x\,{}^p_{x'}] \geqslant \tilde\measuredangle^{\kappa}(x\,{}^p_{x'}) \quad\text{and}\quad \measuredangle[x'\,{}^p_x] \geqslant \tilde\measuredangle^{\kappa}(x'\,{}^p_x).$$

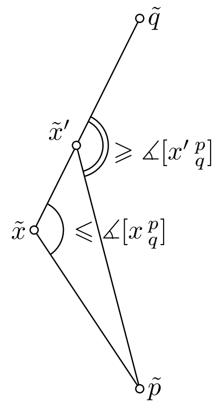

Now, let $[\tilde x\tilde x'\tilde p] = \tilde\triangle^{\kappa}(xx'p)$. Take $\tilde q$ on the extension of $[\tilde x\tilde x']$ beyond x' such that $|\tilde x - \tilde q| = |x - q|$ (and therefore $|\tilde x' - \tilde q| = |x' - q|$). From ❺,

$$\measuredangle[x\,{}^p_q] = \measuredangle[x\,{}^p_{x'}] \geqslant \tilde\measuredangle^{\kappa}(x\,{}^p_{x'}) \quad\Rightarrow\quad \tilde\curlyvee^{\kappa}[x\,{}^q_p] \geqslant |\tilde p - \tilde q|.$$

Hence

$$\begin{aligned}\measuredangle[\tilde x'\,{}^{\tilde p}_{\tilde q}] &= \pi - \tilde\measuredangle^{\kappa}(x'\,{}^p_x) \geqslant \\ &\geqslant \pi - \measuredangle[x'\,{}^p_x] = \\ &= \measuredangle[x'\,{}^p_q],\end{aligned}$$

and ❸ follows.

Let us continue the proof. Set $x_0 = x$. Let us apply inductively the above construction to get a sequence of hinges $[x_n\,{}^p_q]$ with $x_{n+1} = x'_n$. From ❸, we have that the sequence $s_n = \tilde\curlyvee^{\kappa}[x_n\,{}^p_q]$ is nonincreasing.

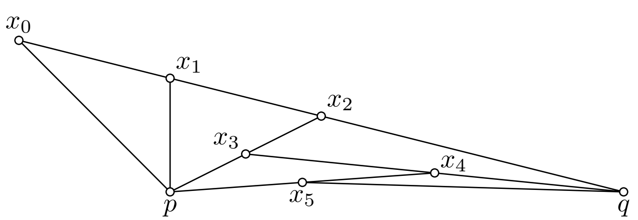

The sequence might terminate at some n only if $|p - x_n| + |x_n - q| < \tfrac{2}{3}\cdot\ell$. In this case, by the assumptions of the lemma, $\tilde\curlyvee^{\kappa}[x_n\,{}^p_q] \geqslant |p - q|$. Since the sequence s_n is nonincreasing, inequality ❶ follows.

Otherwise, the sequence $r_n = |p - x_n| + |x_n - q|$ is nonincreasing and $r_n \geqslant \tfrac{2}{3}\cdot\ell$ for all n. Note that by construction, the distances $|x_n - x_{n+1}|$, $|x_n - p|$, and $|x_n - q|$ are bounded away from zero for all large n. Indeed, since on each step we move x_n toward to the point p or q that is further away, the distances $|x_n - p|$ and $|x_n - q|$ become about the same. Namely,

[2] In fact by 8.37 the condition $[x'q] \subset [xq]$ always holds.

by ❹, we have that $|p-x_n|-|x_n-q|\leqslant\tfrac29\cdot\ell$ for all large n. Since $|p-x_n|+|x_n-q|\geqslant\tfrac23\cdot\ell$, we have $|x_n-p|\geqslant\tfrac{\ell}{100}$ and $|x_n-q|\geqslant\tfrac{\ell}{100}$. Further, since $r_n\geqslant\tfrac23\cdot\ell$, ❹ implies that $|x_n-x_{n+1}|>\tfrac{\ell}{100}$.

Since the sequence r_n is nonincreasing, it converges. In particular $r_n-r_{n+1}\to 0$ as $n\to\infty$. It follows that $\tilde\measuredangle^\kappa(x_n{}^{p_n}_{x_{n+1}})\to\pi$, where $p_n=p$ if $x_{n+1}\in[x_nq]$, and otherwise $p_n=q$. Since $\measuredangle[x_n{}^{p_n}_{x_{n+1}}]\geqslant\tilde\measuredangle^\kappa(x_n{}^{p_n}_{x_{n+1}})$, we have $\measuredangle[x_n{}^{p_n}_{x_{n+1}}]\to\pi$ as $n\to\infty$.

It follows that

$$|p-x_n|+|x_n-q|-\tilde\Upsilon^\kappa[x_n{}^p_q]\to 0.$$

(Here we used that $\ell\leqslant\varpi^\kappa$.) Together with the triangle inequality

$$|p-x_n|+|x_n-q|\geqslant|p-q|$$

this yields

$$\lim_{n\to\infty}\tilde\Upsilon^\kappa[x_n{}^p_q]\geqslant|p-q|.$$

Applying monotonicity of the sequence $s_n=\tilde\Upsilon^\kappa[x_n{}^p_q]$, we obtain ❶. □

The final part of the proof above resembles the cat's cradle construction introduced by the first author and Richard Bishop [3].

The following lemma works in all complete spaces; it will be used as a substitute for the existence of a minimum point of a continuous function on a compact space.

8.35. Lemma on almost minimum. *Let $\mathcal X$ be a complete metric space. Suppose $p\in\mathcal X$, $r\colon\mathcal X\to\mathbb R$ is a function, and $\varepsilon>0$. Assume that the function r is strictly positive in $\overline{\mathrm B}[p,\tfrac{1}{\varepsilon^2}\cdot r(p)]$ and $\underline{\lim}_n\, r(x_n)>0$ for any convergent sequence $x_n\to x\in\overline{\mathrm B}[p,\tfrac{1}{\varepsilon^2}\cdot r(p)]$.*

Then there is a point $p^*\in\overline{\mathrm B}[p,\tfrac{1}{\varepsilon^2}\cdot r(p)]$ such that

a) $r(p^*)\leqslant r(p)$ and

b) $r(x)>(1-\varepsilon)\cdot r(p^*)$ for any $x\in\overline{\mathrm B}[p^*,\tfrac1\varepsilon\cdot r(p^*)]$.

Proof. Assume the statement is wrong. Then for any $x\in\mathrm B(p,\tfrac{1}{\varepsilon^2}\cdot r(p))$ with $r(x)\leqslant r(p)$, there is a point $x'\in\mathcal X$ such that

$$|x-x'|<\tfrac1\varepsilon\cdot r(x)\quad\text{and}\quad r(x')\leqslant(1-\varepsilon)\cdot r(x).$$

Take $x_0=p$ and consider a sequence of points x_n such that $x_{n+1}=x_n'$. Clearly

$$|x_{n+1}-x_n|\leqslant\tfrac{r(p)}{\varepsilon}\cdot(1-\varepsilon)^n\quad\text{and}\quad r(x_n)\leqslant r(p)\cdot(1-\varepsilon)^n.$$

In particular, $|p-x_n|<\tfrac{1}{\varepsilon^2}\cdot r(p)$. Therefore (x_n) is Cauchy, $x_n\to x\in\overline{\mathrm B}[p,\tfrac{1}{\varepsilon^2}\cdot r(p)]$ and $\lim_n r(x_n)=0$, a contradiction. □

Proof of the globalization theorem (8.30). Exactly the same argument as in the proof of Theorem 8.11 shows that $\mathcal{L}$ is G-delta geodesic. By Theorem 8.29-6, for any hinge $[x\,{}^p_y]$ in $\mathcal{L}$ the angle $\measuredangle[x\,{}^p_y]$ is defined and moreover, if $x \in\,]yz[$ then

$$\measuredangle[x\,{}^p_y] + \measuredangle[x\,{}^p_z] \leqslant \pi.$$

Let us denote by $\mathrm{ComRad}(p, \mathcal{L})$ (which stands for comparison radius of $\mathcal{L}$ at p) the maximal value (possibly ∞) such that the comparison

$$\measuredangle[x\,{}^p_y] \geqslant \tilde{\measuredangle}^{\kappa}(x\,{}^p_y)$$

holds for any hinge $[x\,{}^p_y]$ with $|p - x| + |x - y| < \mathrm{ComRad}(p, \mathcal{L})$.

As follows from 8.29-3, $\mathrm{ComRad}(p, \mathcal{L}) > 0$ for any $p \in \mathcal{L}$ and

$$\varliminf_{n\to\infty} \mathrm{ComRad}(p_n, \mathcal{L}) > 0$$

for any converging sequence of points $p_n \to p$. That makes it possible to apply the lemma on almost minimum (8.35) to the function $p \mapsto$ $\mapsto \mathrm{ComRad}(p, \mathcal{L})$.

According to the short hinge lemma (8.28), it is sufficient to show that

❻ $$s_0 = \inf_{p\in\mathcal{L}} \mathrm{ComRad}(p, \mathcal{L}) \geqslant \varpi^{\kappa} \quad \text{for any} \quad p \in \mathcal{L}.$$

We argue by contradiction, assuming that ❻ does not hold.

The rest of the proof is easier for geodesic spaces and easier still for compact spaces. Thus we give three different arguments for each of these cases.

Compact case. Assume $\mathcal{L}$ is compact.

By Theorem 8.29-3, $s_0 > 0$. Take a point $p^* \in \mathcal{L}$ such that $r^* =$ $= \mathrm{ComRad}(p^*, \mathcal{L})$ is sufficiently close to s_0 (p^* such that $s_0 \leqslant r^* <$ $< \min\{\varpi^{\kappa}, \tfrac{3}{2}{\cdot}s_0\}$ will do). Then the key lemma (8.34) applied for p^* and ℓ slightly bigger than r^* (say, such that $r^* < \ell < \min\{\varpi^{\kappa}, \tfrac{3}{2}{\cdot}s_0\}$) implies that

$$\measuredangle[x\,{}^{p^*}_q] \geqslant \tilde{\measuredangle}^{\kappa}(x\,{}^{p^*}_q)$$

for any hinge $[x\,{}^{p^*}_q]$ such that $|p^* - x| + |x - q| < \ell$. Thus $r^* \geqslant \ell$, a contradiction.

Geodesic case. Assume $\mathcal{L}$ is geodesic.

Fix a small $\varepsilon > 0$ ($\varepsilon = 0.0001$ will do). Apply the lemma on almost minimum (8.35) to find a point $p^* \in \mathcal{L}$ such that

$$r^* = \mathrm{ComRad}(p^*, \mathcal{L}) < \varpi^{\kappa}$$

and

$$\text{❼} \qquad \operatorname{ComRad}(q,\mathcal{L}) > (1-\varepsilon)\cdot r^*$$

for any $q \in \overline{\mathrm{B}}[p^*, \tfrac{1}{\varepsilon}\cdot r^*]$.

Applying the key lemma (8.34) for p^* and ℓ slightly bigger than r^* leads to a contradiction.

General case. Let us construct $p^* \in \mathcal{L}$ as in the previous case. Since $\mathcal{L}$ is not geodesic, we cannot apply the key lemma directly. Instead, let us pass to the ultrapower $\mathcal{L}^\omega$, which is a geodesic space (see 3.6).

In Theorem 8.29, inequality ❼ implies that condition 8.29-1 holds for some fixed $R_1 = \frac{r^*}{100} > 0$ at any point $q \in \overline{\mathrm{B}}[p^*, \tfrac{1}{2\cdot\varepsilon}\cdot r^*] \subset \mathcal{L}$. Therefore a similar statement is true in the ultrapower $\mathcal{L}^\omega$; that is, for any point $q_\omega \in \overline{\mathrm{B}}[p^*, \tfrac{1}{2\cdot\varepsilon}\cdot r^*] \subset \mathcal{L}^\omega$, condition 8.29-1 holds for, say, $R_1 = \frac{r^*}{101}$.

Note that $r^* \geqslant \operatorname{ComRad}(p^*, \mathcal{L}^\omega)$. Therefore we can apply the lemma on almost minimum at the point p^* to the function $x \mapsto \operatorname{ComRad}(x, \mathcal{L}^\omega)$ and $\varepsilon' = \sqrt{\varepsilon} = 0.01$.

For the resulting point $p^{**} \in \mathcal{L}^\omega$, we have $r^{**} = \operatorname{ComRad}(p^{**}, \mathcal{L}) < \varpi^\kappa$, and $\operatorname{ComRad}(q_\omega, \mathcal{L}^\omega) > (1-\varepsilon')\cdot r^{**}$ for any $q_\omega \in \overline{\mathrm{B}}[p^{**}, \tfrac{1}{\varepsilon'}\cdot r^{**}]$. Thus applying the key lemma (8.34) for p^{**} and for ℓ slightly bigger than r^{**} leads to a contradiction. □

G Properties of geodesics and angles

Remark. All proofs in this section can be easily modified to use only the local definition of CBB spaces without use of the globalization theorem (8.30).

8.36. Geodesics do not split. *In a* CBB *space, geodesics do not bifurcate.*

More precisely, let $\mathcal{L}$ *be a* CBB *space and* $[px]$, $[py]$ *be two geodesics. Then:*

a) If there is $\varepsilon > 0$ *such that* $\operatorname{geod}_{[px]}(t) = \operatorname{geod}_{[py]}(t)$ *for all* $t \in [0,\varepsilon)$*, then* $[px] \subset [py]$ *or* $[py] \subset [px]$.

b) If $\measuredangle[p\,{}^x_y] = 0$*, then* $[px] \subset [py]$ *or* $[py] \subset [px]$.

8.37. Corollary. *Let* $\mathcal{L}$ *be a* CBB *space. Then the restriction of any geodesic in* $\mathcal{L}$ *to a proper segment is the unique minimal geodesic joining its endpoints.*

In case $\kappa \leqslant 0$, the proof is easier, since the model triangles are always defined. To deal with $\kappa > 0$ we have to argue locally.

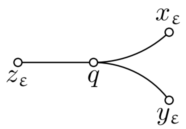

Proof of 8.36; (a). Let $t_{\max}$ be the maximal value such that $\mathrm{geod}_{[px]}(t) = \mathrm{geod}_{[py]}(t)$ for all $t \in [0, t_{\max})$. Since geodesics are continuous, $\mathrm{geod}_{[px]}(t_{\max}) = {}$ $= \mathrm{geod}_{[py]}(t_{\max})$. Let

$$q = \mathrm{geod}_{[px]}(t_{\max}) = \mathrm{geod}_{[py]}(t_{\max}).$$

We must show that $t_{\max} = \min\{|p-x|, |p-y|\}$.

If that is not true, choose a sufficiently small $\varepsilon > 0$ such that the points

$$x_\varepsilon = \mathrm{geod}_{[px]}(t_{\max}+\varepsilon) \quad\text{and}\quad y_\varepsilon = \mathrm{geod}_{[py]}(t_{\max}+\varepsilon)$$

are distinct. Let

$$z_\varepsilon = \mathrm{geod}_{[px]}(t_{\max}-\varepsilon) = \mathrm{geod}_{[py]}(t_{\max}-\varepsilon).$$

Clearly, $\tilde\measuredangle^\kappa(q\,{}^{z_\varepsilon}_{x_\varepsilon}) = \tilde\measuredangle^\kappa(q\,{}^{z_\varepsilon}_{y_\varepsilon}) = \pi$. Thus from the CBB($\kappa$) comparison (8.2), $\tilde\measuredangle^\kappa(q\,{}^{x_\varepsilon}_{y_\varepsilon}) = 0$ and thus $x_\varepsilon = y_\varepsilon$, a contradiction.

(b) From hinge comparison 8.14c,

$$\measuredangle[p\,{}^x_y] = 0 \quad\Rightarrow\quad \tilde\measuredangle^\kappa\left(p\,{}^{\mathrm{geod}_{[px]}(t)}_{\mathrm{geod}_{[py]}(t)}\right) = 0$$

and thus $\mathrm{geod}_{[px]}(t) = \mathrm{geod}_{[py]}(t)$ for all small t. Therefore we can apply (a). □

8.38. Adjacent angle lemma. *Let $\mathcal{L}$ be a* CBB *space. Assume that two hinges $[z\,{}^x_p]$ and $[z\,{}^y_p]$ in $\mathcal{L}$ are* a d j a c e n t*; that is, they share a common side $[zp]$ and $z \in\,]xy[$. Then*

$$\measuredangle[z\,{}^p_y] + \measuredangle[z\,{}^p_x] = \pi.$$

Proof. From the hinge comparison (8.14c) we have that both angles $\measuredangle[z\,{}^p_y]$ and $\measuredangle[z\,{}^p_x]$ are defined and

$$\measuredangle[z\,{}^p_y] + \measuredangle[z\,{}^p_x] \leqslant \pi.$$

Clearly $\measuredangle[z\,{}^x_y] = \pi$. Thus the result follows from the triangle inequality for angles (6.4). □

8.39. Angle semicontinuity. *Suppose $\mathcal{L}_n$ is a sequence of* CBB(κ) *spaces and $\mathcal{L}_n \to \mathcal{L}_\varpi$ as $n \to \varpi$. Assume that a sequence of hinges $[p_n\,{}^{x_n}_{y_n}]$ in $\mathcal{L}_n$ converges to a hinge $[p_\varpi\,{}^{x_\varpi}_{y_\varpi}]$ in $\mathcal{L}_\varpi$. Then*

$$\measuredangle[p_\varpi\,{}^{x_\varpi}_{y_\varpi}] \leqslant \lim_{n\to\varpi} \measuredangle[p_n\,{}^{x_n}_{y_n}].$$

Proof. From 8.17,

$$\measuredangle[p_\varpi\,{}^{x_\varpi}_{y_\varpi}] = \sup\left\{\, \tilde\measuredangle^\kappa(p_\varpi\,{}^{\bar x_\varpi}_{\bar y_\varpi}) : \bar x_\varpi \in\,]p_\varpi x_\varpi],\ \bar y_\varpi \in\,]p_\varpi x_\varpi] \,\right\}.$$

For fixed $\bar x_\omega \in\,]p_\omega x_\omega]$ and $\bar y_\omega \in\,]p_\omega x_\omega]$, choose $\bar x_n \in$ $\in\,]px_n]$ and $\bar y_n \in\,]py_n]$ so that $\bar x_n \to \bar x_\omega$ and $\bar y_n \to \bar y_\omega$ as $n \to \omega$. Clearly

$$\tilde\measuredangle^\kappa(p_n\,{}^{\bar x_n}_{\bar y_n}) \to \tilde\measuredangle^\kappa(p_\omega\,{}^{\bar x_\omega}_{\bar y_\omega})$$

as $n \to \omega$.

From the hinge comparison (8.14c), $\measuredangle[p_n\,{}^{x_n}_{y_n}] \geqslant$ $\geqslant \tilde\measuredangle^\kappa(p_n\,{}^{\bar x_n}_{\bar y_n})$. Hence the result. □

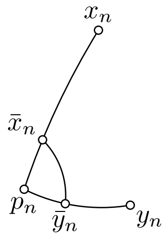

8.40. Angle continuity. *Let $\mathcal{L}_n$ be a sequence of complete length* CBB(κ) *spaces, and $\mathcal{L}_n \to \mathcal{L}_\omega$ as $n \to \omega$. Assume that sequences of points p_n, x_n, y_n in $\mathcal{L}_n$ converge to points $p_\omega, x_\omega, y_\omega$ in $\mathcal{L}_\omega$, and the following two conditions hold:*

a) $p_\omega \in \operatorname{Str}(x_\omega)$,

b) $p_\omega \in \operatorname{Str}(y_\omega)$ *or* $y_\omega \in \operatorname{Str}(p_\omega)$.

Then

$$\measuredangle[p_\omega\,{}^{x_\omega}_{y_\omega}] = \lim_{n\to\omega} \measuredangle[p_n\,{}^{x_n}_{y_n}].$$

Proof. By Corollary 8.32, we may assume that $\kappa \leqslant 0$.

By Plaut's theorem (8.11), the hinge $[p_\omega\,{}^{x_\omega}_{y_\omega}]$ is uniquely defined. Therefore the hinges $[p_n\,{}^{x_n}_{y_n}]$ converge to $[p_\omega\,{}^{x_\omega}_{y_\omega}]$ as $n \to \omega$. Hence by the angle semicontinuity (8.39), we have

$$\measuredangle[p_\omega\,{}^{x_\omega}_{y_\omega}] \leqslant \lim_{n\to\omega} \measuredangle[p_n\,{}^{x_n}_{y_n}].$$

It remains to show that

❶ $$\measuredangle[p_\omega\,{}^{x_\omega}_{y_\omega}] \geqslant \lim_{n\to\omega} \measuredangle[p_n\,{}^{x_n}_{y_n}].$$

Fix $\varepsilon > 0$. Since $p_\omega \in \operatorname{Str}(x_\omega)$, there is a point $q_\omega \in \mathcal{L}_\omega$ such that

$$\tilde\measuredangle^\kappa(p_\omega\,{}^{x_\omega}_{q_\omega}) > \pi - \varepsilon.$$

The hinge comparison (8.14c) implies that

❷ $$\measuredangle[p_\omega\,{}^{x_\omega}_{q_\omega}] > \pi - \varepsilon.$$

By the triangle inequality for angles (6.4),

❸ $$\begin{aligned}\measuredangle[p_\omega\,{}^{x_\omega}_{y_\omega}] &\geqslant \measuredangle[p_\omega\,{}^{x_\omega}_{q_\omega}] - \measuredangle[p_\omega\,{}^{y_\omega}_{q_\omega}] > \\ &> \pi - \varepsilon - \measuredangle[p_\omega\,{}^{y_\omega}_{q_\omega}].\end{aligned}$$

Note that we can assume in addition that $q_\omega \in \mathrm{Str}(p_\omega)$. Choose $q_n \in \mathcal{L}_n$ such that $q_n \to q_\omega$ as $n \to \omega$. Note that by angle semicontinuity we again have

❹
$$\begin{aligned}\measuredangle[p_\omega\,{}^{x_\omega}_{q_\omega}] &\leqslant \lim_{n\to\omega} \measuredangle[p_n\,{}^{x_n}_{q_n}],\\ \measuredangle[p_\omega\,{}^{y_\omega}_{q_\omega}] &\leqslant \lim_{n\to\omega} \measuredangle[p_n\,{}^{y_n}_{q_n}].\end{aligned}$$

By the CBB(κ) comparison (8.2) and 8.17b,

$$\measuredangle[p_n\,{}^{x_n}_{y_n}] + \measuredangle[p_n\,{}^{y_n}_{q_n}] + \measuredangle[p_n\,{}^{x_n}_{q_n}] \leqslant 2\cdot\pi$$

for all n. Together with ❹, ❷ and ❸, this implies

$$\begin{aligned}\lim_{n\to\omega} \measuredangle[p_n\,{}^{x_n}_{y_n}] &\leqslant 2\cdot\pi - \lim_{n\to\omega} \measuredangle[p_n\,{}^{x_n}_{q_n}] - \lim_{n\to\omega} \measuredangle[p_n\,{}^{y_n}_{q_n}] \leqslant \\ &\leqslant 2\cdot\pi - \measuredangle[p_\omega\,{}^{x_\omega}_{q_\omega}] - \measuredangle[p_\omega\,{}^{y_\omega}_{q_\omega}] < \\ &< \measuredangle[p_\omega\,{}^{x_\omega}_{y_\omega}] + 2\cdot\varepsilon.\end{aligned}$$

Since $\varepsilon > 0$ is arbitrary, ❶ follows. □

8.41. First variation formula. *Let $\mathcal{L}$ be a complete length* CBB *space. For any point q and any geodesic $[px]$ in $\mathcal{L}$ with $p \neq q$, we have*

❺
$$|q - \mathrm{geod}_{[px]}(t)| = |q - p| - t\cdot\cos\varphi + o(t),$$

where φ is the infimum of angles between $[px]$ and all geodesics from p to q in the ultrapower $\mathcal{L}^\omega$.

Remark. If $\mathcal{L}$ is a proper space, then $\mathcal{L}^\omega = \mathcal{L}$, see Section 3B. Therefore the infimum φ is achieved on some particular geodesic from p to q.

As a corollary we obtain the following classical result:

8.42. Strong angle lemma. *Let $\mathcal{L}$ be a complete length* CBB *space and $p \neq q \in \mathcal{L}$ be such that there is unique geodesic from p to q in the ultrapower $\mathcal{L}^\omega$. Then for any hinge $[p\,{}^q_x]$ we have*

❻
$$\measuredangle[p\,{}^q_x] = \lim_{\substack{\bar{x}\to p\\ \bar{x}\in\,]px]}} \tilde{\measuredangle}^\kappa(p\,{}^q_{\bar{x}})$$

for any $\kappa \in \mathbb{R}$ such that $|p - q| < \varpi^\kappa$.

In particular, ❻ *holds if $p \in \mathrm{Str}(q)$ as well as if $q \in \mathrm{Str}(p)$.*

Remark.

- The above lemma is essentially due to Alexandrov. The right-hand side in ❻ is called the s t r o n g a n g l e of the hinge $[p\,{}^q_x]$. Note that in a general metric space the angle and the strong angle of the same hinge might differ.

◇ As follows from Corollary 3.8, if there is a unique geodesic $[pq]$ in the ultrapower $\mathcal{L}^\omega$, then $[pq]$ lies in $\mathcal{L}$.

Proof of 8.42. The first statement follows directly from the first variation formula (8.41) and the definition of model angle (see Section 6C). The second statement follows from Plaut's theorem (8.11) applied to $\mathcal{L}^\omega$. (Note that according to Proposition 8.5, $\mathcal{L}^\omega$ is a complete length CBB space.) □

Proof of 8.41. By Corollary 8.32, we can assume that $\kappa \leqslant 0$. The inequality

$$|q - \mathrm{geod}_{[px]}(t)| \leqslant |q - p| - t\cdot \cos\varphi + o(t)$$

follows from the first variation inequality (6.6). Thus, it is sufficient to show that

$$|q - \mathrm{geod}_{[px]}(t)| \geqslant |q - p| - t\cdot \cos\varphi + o(t).$$

Assume the contrary. Then there is $\varepsilon > 0$ such that $\varphi + \varepsilon < \pi$, and for a sequence $t_n \to 0+$ we have

❼ $$|q - \mathrm{geod}_{[px]}(t_n)| < |q - p| - t_n\cdot \cos(\varphi - \varepsilon).$$

Let $x_n = \mathrm{geod}_{[px]}(t_n)$. Clearly

$$\tilde\measuredangle^\kappa(x_n\,{}^p_q) > \pi - \varphi + \tfrac{\varepsilon}{2}$$

for all large n.

Assume $\mathcal{L}$ is geodesic. Choose a sequence of geodesics $[x_n q]$. Let $[x_n q] \to [pq]_{\mathcal{L}^\omega}$ as $n \to \omega$ (in general $[pq]$ might lie in $\mathcal{L}^\omega$). Applying both parts of hinge comparison (8.14c), we have $\measuredangle[x_n\,{}^q_x] < \varphi - \frac{\varepsilon}{2}$ for all large n. According to 8.39, the angle between $[pq]$ and $[px]$ is at most $\varphi - \frac{\varepsilon}{2}$, a contradiction.

Finally, if $\mathcal{L}$ is not geodesic, choose a sequence $q_n \in \mathrm{Str}(x_n)$, such that $q_n \to q$ and the inequality

$$\tilde\measuredangle^\kappa(x_n\,{}^{p}_{q_n}) > \pi - \varphi + \tfrac{\varepsilon}{2}$$

still holds. Then the same argument as above shows that $[x_n q_n]$ ω-converges to a geodesic $[pq]_{\mathcal{L}^\omega}$ from p to q having angle at most $\varphi - \frac{\varepsilon}{2}$ with $[px]$. □

H On positive lower bound

In this section we consider CBB(κ) spaces for $\kappa > 0$. Applying rescaling we can assume that $\kappa = 1$.

The following theorem states that if one ignores a few exceptional spaces, then the diameter of a space with positive lower curvature bound is bounded. Note that many authors (but not us) exclude these spaces in the definition of Alexandrov space with positive lower curvature bound.

8.43. On diameter of a space. *Let $\mathcal{L}$ be a complete length* CBB(1) *space. Then either*

a) $\operatorname{diam}\mathcal{L} \leqslant \pi$*;*

b) or $\mathcal{L}$ is isometric to one of the following exceptional spaces:

1. *real line $\mathbb{R}$,*
2. *a half-line $\mathbb{R}_{\geqslant 0}$,*
3. *a closed interval $[0,a] \in \mathbb{R}$, $a > \pi$,*
4. *a circle $\mathbb{S}^1_a$ of length $a > 2\cdot\pi$.*

Proof. Assume that $\mathcal{L}$ is a geodesic space and $\operatorname{diam}\mathcal{L} > \pi$. Choose $x, y \in \mathcal{L}$ so that $|x-y| = \pi + \varepsilon$, $0 < \varepsilon < \frac{\pi}{4}$. By moving y slightly, we can also assume that the geodesic $[xy]$ is unique; to prove this, use either Plaut's theorem (8.11) or the fact that geodesics do not split (8.36). Let z be the midpoint of the geodesic $[xy]$.

Consider the function $f = \operatorname{dist}_x + \operatorname{dist}_y$. As follows from Lemma 8.44, f is concave in $\mathrm{B}(z, \frac{\varepsilon}{4})$. Let $p \in \mathrm{B}(z, \frac{\varepsilon}{4})$. Choose a geodesic $[zp]$, and let $h(t) = f \circ \operatorname{geod}_{[zp]}(t)$ and $\ell = |z-p|$. Clearly h is concave. From the adjacent angle lemma (8.38), we have $h^+(0) = 0$. Therefore h is nonincreasing which means that

$$|x-p| + |y-p| = h(\ell) \leqslant h(0) = |x-y|.$$

Since the geodesic $[xy]$ is unique this means that $p \in [xy]$, and hence $\mathrm{B}(z, \frac{\varepsilon}{4})$ only contains points of $[xy]$. Since in CBB spaces, geodesics do not bifurcate (8.36*a*), it follows that all of $\mathcal{L}$ coincides with the maximal extension of $[xy]$ as a local geodesic γ (which might not be minimizing). In other words, $\mathcal{L}$ is isometric to a 1-dimensional Riemannian manifold with possibly nonempty boundary. From this, it is easy to see that $\mathcal{L}$ falls into one of the exceptional spaces described in the theorem.

Lastly, if $\mathcal{L}$ is not geodesic and $\operatorname{diam}\mathcal{L} > \pi$, then the above argument applied to $\mathcal{L}^\omega$ yields that $\mathcal{L}^\omega$ is isometric to one of the exceptional spaces. As all of those spaces are proper, $\mathcal{L} = \mathcal{L}^\omega$. □

8.44. Lemma. *Let $\mathcal{L}$ be a complete length* CBB(1) *space and $p \in \mathcal{L}$. Then $\operatorname{dist}_p : \mathcal{L} \to \mathbb{R}$ is concave in $\mathrm{B}(p,\pi)\backslash\mathrm{B}(p,\frac{\pi}{2})$.*

In particular, if $\operatorname{diam}\mathcal{L} \leqslant \pi$ then the complements $\mathcal{L}\backslash\mathrm{B}(p,r)$ and $\mathcal{L}\backslash\overline{\mathrm{B}}[p,r]$ are convex for any $r > \frac{\pi}{2}$.

Proof. This is a consequence of 8.23*b*. □

8.45. Exercise. *Let $\mathcal{L}$ be an m-dimensional complete length* CBB(1) *space and* $\operatorname{diam}\mathcal{L} \leqslant \pi$. *Assume that a group G acts on $\mathcal{L}$ by isometries, has closed orbits, and*

$$\operatorname{diam}(\mathcal{L}/G) > \tfrac{\pi}{2}.$$

Show that the action of G has a fixed point in $\mathcal{L}$.

8.46. Advanced exercise. *Let $\mathcal{L}$ be a complete length* CBB(1) *space Show that $\mathcal{L}$ contains at most 3 points with space of directions $\leqslant \frac{1}{2}\cdot\mathbb{S}^n$ (see Definition 4.9).*

8.47. On perimeter of a triple. *Suppose $\mathcal{L}$ is a complete length* CBB(1) *space and* $\operatorname{diam}\mathcal{L} \leqslant \pi$. *Then the perimeter of any triple of points $p, q, r \in \mathcal{L}$ is at most $2\cdot\pi$.*

Proof. Arguing by contradiction, suppose

❶ $$|p-q| + |q-r| + |r-p| > 2\cdot\pi$$

for some $p, q, r \in \mathcal{L}$. Rescaling the space slightly, we can assume that $\operatorname{diam}\mathcal{L} < \pi$, but the inequality ❶ still holds.

By Corollary 8.32, after rescaling $\mathcal{L}$ is still CBB(1).

Since $\mathcal{L}$ is G-delta geodesic (8.11), it is sufficient to consider the case when there is a geodesic $[qr]$.

First note that since $\operatorname{diam}\mathcal{L} < \pi$, by 8.23*b* the function

$$y(t) = \operatorname{md}^1|p - \operatorname{geod}_{[qr]}(t)|$$

satisfies the differential inequality $y'' \leqslant 1 - y$.

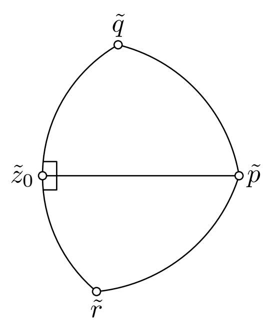

Take $z_0 \in [qr]$ so that the restriction $\operatorname{dist}_p|_{[qr]}$ attains its maximum at z_0, and set $t_0 = |q - z_0|$ so $z_0 = \operatorname{geod}_{[qr]}(t_0)$. Consider the following model configuration: two geodesics $[\tilde p\tilde z_0]$, $[\tilde q\tilde r]$ in $\mathbb{S}^2$ such that

$$|\tilde p - \tilde z_0| = |p - z_0|, \qquad |\tilde q - \tilde r| = |q - r|,$$
$$|\tilde z_0 - \tilde q| = |z_0 - q|, \qquad |\tilde z_0 - \tilde r| = |z_0 - q|$$

and

$$\measuredangle[\tilde z_0\,{}^{\tilde q}_{\tilde p}] = \measuredangle[\tilde z_0\,{}^{\tilde r}_{\tilde p}] = \tfrac{\pi}{2}.$$

Clearly, $\bar y(t) = \operatorname{md}^1|\tilde p - \operatorname{geod}_{[\tilde q\tilde r]}(t)|$ satisfies $\bar y'' = 1 - \bar y$ and $\bar y'(t_0) = 0$, $\bar y(t_0) = y(t_0)$. Since z_0 is a maximum point, $y(t) \leqslant y(t_0) + o(t - t_0)$; thus,

$\bar{y}(t)$ is a barrier for $y(t)=\mathrm{md}^1|p-\mathrm{geod}_{[qr]}(t)|$ at t_0 by 5.14d. From the barrier inequality 5.14d, we get

$$|\tilde p-\mathrm{geod}_{[\tilde q\tilde r]}(t)|\geqslant|p-\mathrm{geod}_{[qr]}(t)|,$$

and hence $|\tilde p-\tilde q|\geqslant|p-q|$ and $|\tilde p-\tilde r|\geqslant|p-r|$.

Therefore $|p-q|+|q-r|+|r-p|$ cannot exceed the perimeter of the spherical triangle $[\tilde p\tilde q\tilde r]$; that is,

$$|p-q|+|q-r|+|r-p|\leqslant 2\cdot\pi,$$

a contradiction. ☐

Let $\kappa>0$. Consider the following extension $\tilde\measuredangle^{\kappa+}(*\,{}^*_*)$ of the model angle function $\tilde\measuredangle^{\kappa}(*\,{}^*_*)$. This definition works well for CBB spaces; for CAT spaces there is a similar but different definition. Some authors define the comparison angle to be $\tilde\measuredangle^{\kappa+}(*\,{}^*_*)$.

8.48. Definition of extended angle. *Suppose p,q,r are points in a metric space, and $p\ne q$, $p\ne r$. Let*

$$\tilde\measuredangle^{\kappa+}(p\,{}^q_r)=\sup\left\{\,\tilde\measuredangle^{\mathrm K}(p\,{}^q_r)\;:\;\mathrm K\leqslant\kappa\,\right\}.$$

The value $\tilde\measuredangle^{\kappa+}(p\,{}^q_r)$ is called the extended model angle of the triple p, q, r.

8.49. Extended angle comparison. *Let $\kappa>0$ and $\mathcal L$ be a complete length* $\mathrm{CBB}(\kappa)$ *space. Then for any hinge $[p\,{}^q_r]$ we have $\measuredangle[p\,{}^q_r]\geqslant\tilde\measuredangle^{\kappa+}(p\,{}^q_r)$.*

Moreover, the extended model angle $\tilde\measuredangle^{\kappa+}(p\,{}^q_r)$ can be calculated using the following rule:

a) $\tilde\measuredangle^{\kappa+}(p\,{}^q_r)=\tilde\measuredangle^{\kappa}(p\,{}^q_r)$ if $\tilde\measuredangle^{\kappa}(p\,{}^q_r)$ is defined;

b) $\tilde\measuredangle^{\kappa+}(p\,{}^q_r)=\tilde\measuredangle^{\kappa+}(p\,{}^r_q)=0$ if $|p-q|+|q-r|=|p-r|$;

c) $\tilde\measuredangle^{\kappa+}(p\,{}^q_r)=\pi$ if none of the above is applicable.

Proof. From Corollary 8.32, $\mathrm K<\kappa$ implies that any complete length $\mathrm{CBB}(\mathrm K)$ space is $\mathrm{CBB}(\kappa)$; thus the extended angle comparison follows from the definition.

The rule for calculating extended angle is an easy consequence of its definition. ☐

I Remarks and open problems

The question whether the first part of 8.14c suffices to conclude that $\mathcal L$ is $\mathrm{CBB}(\kappa)$ is a long-standing open problem (possibly dating back to

Alexandrov), but as far as we know it was first stated in print in [27, footnote in 4.1.5].

8.50. Open question. *Let $\mathcal{L}$ be a complete geodesic space (you can also assume that $\mathcal{L}$ is homeomorphic to $\mathbb{S}^2$ or $\mathbb{R}^2$) such that for any hinge $[x\,{}^p_y]$ in $\mathcal{L}$, the angle $\measuredangle[x\,{}^p_y]$ is defined and*

$$\measuredangle[x\,{}^p_y] \geqslant \tilde\measuredangle^0(x\,{}^p_y).$$

Is it true that $\mathcal{L}$ is CBB(0)*?*

Chapter 9

Fundamentals of curvature bounded above

A Four-point comparison.

9.1. Four-point comparison. *A quadruple of points* p^1, p^2, x^1, x^2 *in a metric space satisfies* CAT(κ) *comparison if*

a) $\tilde\measuredangle^\kappa(p^1\,{}^{x^1}_{x^2}) \leqslant \tilde\measuredangle^\kappa(p^1\,{}^{p^2}_{x^1}) + \tilde\measuredangle^\kappa(p^1\,{}^{p^2}_{x^2})$, *or*

b) $\tilde\measuredangle^\kappa(p^2\,{}^{x^1}_{x^2}) \leqslant \tilde\measuredangle^\kappa(p^2\,{}^{p^1}_{x^1}) + \tilde\measuredangle^\kappa(p^2\,{}^{p^1}_{x^2})$, *or*

c) *one of the six model angles*

$$\tilde\measuredangle^\kappa(p^1\,{}^{x^1}_{x^2}),\quad \tilde\measuredangle^\kappa(p^1\,{}^{p^2}_{x^1}),\quad \tilde\measuredangle^\kappa(p^1\,{}^{p^2}_{x^2}),$$
$$\tilde\measuredangle^\kappa(p^2\,{}^{x^1}_{x^2}),\quad \tilde\measuredangle^\kappa(p^2\,{}^{p^1}_{x^1}),\quad \tilde\measuredangle^\kappa(p^2\,{}^{p^1}_{x^2})$$

is undefined.

Here is a more intuitive pointwise formulation.

9.2. Reformulation. *Let* $\mathcal{X}$ *be a metric space. A quadruple* $p^1, p^2, x^1, x^2 \in \mathcal{X}$ *satisfies* CAT(κ) *comparison if one of the following holds:*

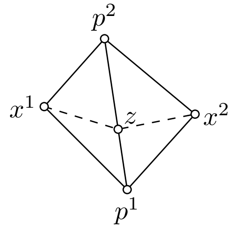

a) *One of the triples* (p^1, p^2, x^1) *or* (p^1, p^2, x^2) *has perimeter* $> 2\cdot\varpi^\kappa$.

b) *If* $[\tilde p^1\tilde p^2\tilde x^1] = \tilde\triangle^\kappa(p^1p^2x^1)$ *and* $[\tilde p\tilde p^2\tilde x^2] = \tilde\triangle^\kappa p^1p^2x^2$, *then*

$$|\tilde x^1 - \tilde z| + |\tilde z - \tilde x^2| \geqslant |x^1 - x^2|,$$

for any $\tilde z \in [\tilde p^1\tilde p^2]$.

9.3. Definition. *Let $\mathcal{U}$ be a metric space.*

a) $\mathcal{U}$ is CAT(κ) *if any quadruple $p^1, p^2, x^1, x^2 \in \mathcal{X}$ satisfies* CAT(κ) *comparison.*

b) $\mathcal{U}$ is locally CAT(κ) *if any point $q \in \mathcal{U}$ admits a neighborhood $\Omega \ni q$ such that any quadruple $p^1, p^2, x^1, x^2 \in \mathcal{X}$ satisfies* CAT(κ) *comparison.*

c) $\mathcal{U}$ is a CAT *space if $\mathcal{U}$ is* CAT(κ) *for some $\kappa \in \mathbb{R}$.*

The condition $\mathcal{U}$ is CAT(κ) should be understood as "$\mathcal{U}$ has global curvature $\leqslant \kappa$". In Proposition 9.18, it will be shown that this formulation makes sense; in particular, if $\kappa \leqslant K$, then any CAT(κ) space is CAT(K).

This terminology was introduced by Michael Gromov; CAT stands for Élie Cartan, Alexandr Alexandrov, and Victor Toponogov. Originally these spaces were called $\mathfrak{R}_\kappa$ domains; this is Alexandrov's terminology and is still in use.

9.4. Exercise. *Let $\mathcal{U}$ be a metric space. Show that $\mathcal{U}$ is* CAT(κ) *if and only if every quadruple of points in $\mathcal{U}$ admits a labeling by (p, x^1, x^2, x^3) such that the three angles $\tilde\measuredangle^\kappa(p\,{}^{x^1}_{x^2})$, $\tilde\measuredangle^\kappa(p\,{}^{x^2}_{x^3})$ and $\tilde\measuredangle^\kappa(p\,{}^{x^1}_{x^3})$ satisfy all three triangle inequalities or one of these angles is undefined. (This is the CAT analog of Definition 8.2 of CBB.)*

9.5. Exercise. *Show that $\mathcal{U}$ is* CAT(κ) *if and only if for any quadruple of points p^1, p^2, x^1, x^2 in $\mathcal{U}$ such that $|p^1-p^2|, |x^1-x^2| \leqslant \varpi^\kappa$, there is a quadruple q^1, q^2, y^1, y^2 in $\mathbb{M}^m(\kappa)$ such that*

$$|q^1-q^2| = |p^1-p^2|, \quad |y^1-y^2| = |x^1-x^2|, \quad |q^i-y^j| \leqslant |p^i-x^j|$$

for any i and j.

9.6. Advanced exercise. *Let $\mathcal{U}$ be a complete length space such that for any quadruple $p, q, x, y \in \mathcal{L}$ the following inequality holds*

$$❶\qquad |p-q|^2 + |x-y|^2 \leqslant |p-x|^2 + |p-y|^2 + |q-x|^2 + |q-y|^2.$$

Prove that $\mathcal{U}$ is CAT(0).

Construct a 4-point metric space $\mathcal{X}$ that satisfies inequality ❶ for any relabeling of its points by p, q, x, y, and such that $\mathcal{X}$ is not CAT(0).

The next proposition follows directly from Definition 9.3 and the definitions of ultralimit and ultrapower; see Section 3B for the related definitions. Recall that ω denotes a fixed selective ultrafilter on $\mathbb{N}$.

9.7. Proposition. *Let $\mathcal{U}_n$ be a* CAT(κ_n) *space for each $n \in \mathbb{N}$. Assume $\mathcal{U}_n \to \mathcal{U}_\omega$ and $\kappa_n \to \kappa_\omega$ as $n \to \omega$. Then $\mathcal{U}_\omega$ is* CAT(κ_ω).

Moreover, a metric space $\mathcal{U}$ is CAT(κ) *if and only if so is its ultrapower $\mathcal{U}^\omega$.*

B Geodesics

9.8. Uniqueness of geodesics. *In a complete length* CAT(κ) *space, pairs of points at distance* $< \varpi^\kappa$ *are joined by unique geodesics, and these geodesics depend continuously on their endpoint pairs.*

Proof. Fix a complete length CAT(κ) space $\mathcal{U}$. Fix two points $p^1, p^2 \in \mathcal{U}$ such that

$$|p^1 - p^2|_{\mathcal{U}} < \varpi^\kappa.$$

Choose a sequence of approximate midpoints z_n between p^1 and p^2; that is,

❶ $$|p^1 - z_n|, |p^2 - z_n| \to \tfrac{1}{2}\cdot|p^1 - p^2| \quad \text{as} \quad n \to \infty.$$

By the law of cosines, $\tilde\measuredangle^\kappa(p^1\,{}^{z_n}_{p^2})$ and $\tilde\measuredangle^\kappa(p^2\,{}^{z_n}_{p^1})$ are arbitrarily small when n is sufficiently large.

Let us apply CAT(κ) comparison (9.1) to the quadruple p^1, p^2, z_n, $z_{\bar k}$ with large n and $\bar k$. We conclude that $\tilde\measuredangle^\kappa(p\,{}^{z_n}_{z_{\bar k}})$ is arbitrarily small when $n, \bar k$ are sufficiently large and p is either p^1 or p^2. By ❶ and the law of cosines, (z_n) converges.

Since $\mathcal{U}$ is complete, the sequence z_n converges to a midpoint between p^1 and p^2. By Lemma 2.12 we obtain the existence of a geodesic $[p^1p^2]$.

Now suppose $p^1_n \to p^1$, $p^2_n \to p^2$ as $n \to \infty$. Let z_n be the midpoint of a geodesic $[p^1_n p^2_n]$ and z be the midpoint of a geodesic $[p^1p^2]$.

It suffices to show that

❷ $$|z_n - z| \to 0 \quad \text{as} \quad n \to \infty.$$

By the triangle inequality, the z_n are approximate midpoints between p^1 and p^2. Apply the CAT(κ) comparison (9.1) to the quadruple p^1, p^2, z_n, z. For $p = p^1$ or $p = p^2$, we see that $\tilde\measuredangle^\kappa(p\,{}^{z_n}_{z})$ is arbitrarily small when n is sufficiently large. By the law of cosines, ❷ follows. □

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

Moreover the same holds for any locally geodesic locally CAT *space* $\mathcal{U}$ *with nontrivial local homology groups at any point; the latter holds in particular if* $\mathcal{U}$ *is a homological manifold.*

9.10. Exercise. *Assume* $\mathcal{U}$ *is a locally compact geodesic* CAT *space with extendable geodesics; that is, any geodesic in* $\mathcal{U}$ *can be extended to a both-sided infinite local geodesic.*

Show that the space of geodesic directions Σ'_p is complete for any $p \in \in \mathcal{U}$.

By the uniqueness of geodesics (9.8), we have the following.

9.11. Corollary. *Any complete length* $\mathrm{CAT}(\kappa)$ *space is* ϖ^κ*-geodesic.*

9.12. Proposition. *The completion of any geodesic* $\mathrm{CAT}(\kappa)$ *space is a complete length* $\mathrm{CAT}(\kappa)$ *space.*

Moreover, $\mathcal{U}$ *is a geodesic* $\mathrm{CAT}(\kappa)$ *space if and only if there is a complete length* $\mathrm{CAT}(\kappa)$ *space* $\bar{\mathcal{U}}$ *that contains a* ϖ^κ*-convex dense set isometric to* $\mathcal{U}$*.*

Proof. By Theorem 9.8, in order to show that the completion $\bar{\mathcal{U}}$ of any geodesic $\mathrm{CAT}(\kappa)$ space $\mathcal{U}$ is $\mathrm{CAT}(\kappa)$, it is sufficient to verify that the completion of a length space is a length space; this is straightforward.

For the second part of the proposition, note that the completion $\bar{\mathcal{U}}$ contains the original space $\mathcal{U}$ as a dense ϖ^κ-convex subset, and the metric on $\mathcal{U}$ coincides with the induced length metric from $\bar{\mathcal{U}}$. □

Here is a corollary from Proposition 9.12 and Theorem 9.8.

9.13. Corollary. *Let* $\mathcal{U}$ *be a* ϖ^κ*-geodesic* $\mathrm{CAT}(\kappa)$ *space. Then pairs of points in* $\mathcal{U}$ *at distance less than* ϖ^κ *are joined by unique geodesics, and these geodesics depend continuously on their endpoint pairs.*

Moreover for any pair of points $p, q \in \mathcal{U}$ *and any value*

$$\ell > \sup\left\{ \frac{\mathrm{sn}^\kappa r}{\mathrm{sn}^\kappa |p-q|} \;:\; 0 \leqslant r \leqslant |p-q| \right\}$$

there are neighborhoods $\Omega_p \ni p$ *and* $\Omega_q \ni q$ *such that the map*

$$(x, y, t) \mapsto \mathrm{path}_{[xy]}(t)$$

is ℓ*-Lipschitz in* $\Omega_p \times \Omega_q \times [0,1]$*.*

Proof. By Proposition 9.12 any geodesic $\mathrm{CAT}(\kappa)$ space is isometric to a convex dense subset of a complete length $\mathrm{CAT}(\kappa)$ space. It remains to apply Theorem 9.8. □

C More comparisons

Here we give a few reformulations of Definition 9.3.

9.14. Theorem. *If $\mathcal{U}$ is a* CAT(κ) *space, then the following conditions hold for all triples $p, x, y \in \mathcal{U}$ of perimeter $< 2\cdot\varpi^\kappa$:*

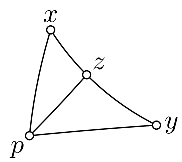

a) (adjacent angle comparison) for any geodesic $[xy]$ and $z \in \,]xy[$, we have

$$\tilde\measuredangle^\kappa(z\,{}^p_x) + \tilde\measuredangle^\kappa(z\,{}^p_y) \geqslant \pi.$$

b) (point-on-side comparison) for any geodesic $[xy]$ and $z \in \,]xy[$, we have

$$\tilde\measuredangle^\kappa(x\,{}^p_y) \geqslant \tilde\measuredangle^\kappa(x\,{}^p_z),$$

or equivalently,

$$|\tilde p - \tilde z| \geqslant |p - z|,$$

where $[\tilde p\tilde x\tilde y] = \tilde\triangle^\kappa(pxy)$, $\tilde z \in \,]\tilde x\tilde y[$, $|\tilde x - \tilde z| = |x - z|$.

c) (hinge comparison) for any hinge $[x\,{}^p_y]$, the angle $\measuredangle[x\,{}^p_y]$ exists and

$$\measuredangle[x\,{}^p_y] \leqslant \tilde\measuredangle^\kappa(x\,{}^p_y),$$

or equivalently,

$$\tilde\Upsilon^\kappa[x\,{}^p_y] \leqslant |p - y|.$$

Moreover, if $\mathcal{U}$ is ϖ^κ-geodesic, then the converse holds in each case.

Remark. In the following proof, the part $(c)\Rightarrow(a)$ only requires that the CAT(κ) comparison (9.1) hold for any quadruple, and does not require the existence of geodesics at distance $< \varpi^\kappa$. The same is true of the parts $(a)\Leftrightarrow(b)$ and $(b)\Rightarrow(c)$. Thus the conditions (a), (b) and (c) are valid for any metric space (not necessarily a length space) that satisfies CAT(κ) comparison (9.1). The converse does not hold; for example, all these conditions are vacuously true in a totally disconnected space, while CAT(κ) comparison is not.

Proof; (a). Since the perimeter of p, x, y is $< 2\cdot\varpi^\kappa$, so is the perimeter of any subtriple of p, z, x, y by the triangle inequality. By Alexandrov's lemma (6.2),

$$\tilde\measuredangle^\kappa(p\,{}^z_x) + \tilde\measuredangle^\kappa(p\,{}^z_y) < \tilde\measuredangle^\kappa(p\,{}^x_y) \quad \text{or} \quad \tilde\measuredangle^\kappa(z\,{}^p_x) + \tilde\measuredangle^\kappa(z\,{}^p_y) = \pi.$$

In the former case, the CAT(κ) comparison (9.1) applied to the quadruple p, z, x, y implies

$$\tilde\measuredangle^\kappa(z\,{}^p_x) + \tilde\measuredangle^\kappa(z\,{}^p_y) \geqslant \tilde\measuredangle^\kappa(z\,{}^x_y) = \pi.$$

(a) $\Leftrightarrow$ (b). Follows from Alexandrov's lemma (6.2).

(b) ⇒ (c). By (b), for $\bar p \in\,]xp]$ and $\bar y \in\,]xy]$ the function $(|x-\bar p|, |x-\bar y|) \mapsto \tilde\measuredangle^\kappa(x\,{}^{\bar p}_{\bar y})$ is nondecreasing in each argument. In particular, $\measuredangle[x\,{}^{p}_{y}] = \inf \tilde\measuredangle^\kappa(x\,{}^{\bar p}_{\bar y})$. Thus $\measuredangle[x\,{}^{p}_{y}]$ exists and is at most $\tilde\measuredangle^\kappa(x\,{}^{p}_{y})$.

Converse. Assume $\mathcal U$ is ϖ^κ-geodesic. Let us first show that in this case $(c) \Rightarrow (a)$.

Indeed, by (c) and the triangle inequality for angles (6.4),

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

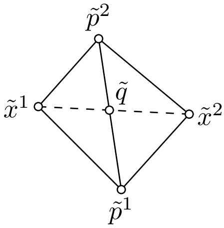

Therefore, it is sufficient to prove the converse for *(b)*.

Given a quadruple p^1, p^2, x^1, x^2 whose subtriples have perimeter $< 2\cdot\varpi^\kappa$, we must verify the CAT(κ) comparison (9.1). In $\mathbb M^2(\kappa)$, construct the model triangles $[\tilde p^1\tilde p^2\tilde x^1] = \tilde\triangle^\kappa(p^1p^2x^1)$ and $[\tilde p^1\tilde p^2\tilde x^2] = \tilde\triangle^\kappa(p^1p^2x^2)$, lying on either side of a common segment $[\tilde p^1\tilde p^2]$. We may suppose

$$\tilde\measuredangle^\kappa(p^1\,{}^{p^2}_{x^1}) + \tilde\measuredangle^\kappa(p^1\,{}^{p^2}_{x^2}) \leqslant \pi \quad\text{and}\quad \tilde\measuredangle^\kappa(p^2\,{}^{p^1}_{x^1}) + \tilde\measuredangle^\kappa(p^2\,{}^{p^1}_{x^2}) \leqslant \pi,$$

since otherwise CAT(κ) comparison holds trivially. Then $[\tilde p^1\tilde p^2]$ and $[\tilde x^1\tilde x^2]$ intersect, say at $\tilde q$.

By assumption, there is a geodesic $[p^1p^2]$. Choose $q \in [p^1p^2]$ corresponding to $\tilde q$; that is, $|p^1-q| = |\tilde p^1-\tilde q|$. Then

$$|x^1-x^2| \leqslant |x^1-q| + |q-x^2| \leqslant |\tilde x^1-\tilde q| + |\tilde q-\tilde x^2| = |\tilde x^1-\tilde x^2|,$$

where the second inequality follows from (b). Therefore by monotonicity of the function $a \mapsto \tilde\measuredangle^\kappa\{a;b,c\}$ (1.1c),

$$\tilde\measuredangle^\kappa(p^1\,{}^{x^1}_{x^2}) \leqslant \measuredangle[\tilde p^1\,{}^{\tilde x^1}_{\tilde x^2}] = \tilde\measuredangle^\kappa(p^1\,{}^{p^2}_{x^1}) + \tilde\measuredangle^\kappa(p^1\,{}^{p^2}_{x^2}).$$

□

Let us display a corollary of the proof of 9.14, namely, monotonicity of the model angle with respect to adjacent sidelengths.

9.15. Angle-sidelength monotonicity. *Suppose $\mathcal U$ is a ϖ^κ-geodesic* CAT(κ) *space, and $p, x, y \in \mathcal U$ have perimeter $< 2\cdot\varpi^\kappa$. Then for $\bar y \in\,]xy]$, the function*

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

is nondecreasing.

In particular, if $\bar p \in\,]xp]$, then

a) the function

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

is nondecreasing in each argument,

b) the angle $\measuredangle[x\,{}^p_y]$ *exists and*

$$\measuredangle[x\,{}^p_y] = \inf\left\{\,\tilde\measuredangle^\kappa(x\,{}^{\bar p}_{\bar y}) : \bar p \in\,]xp],\ \bar y \in\,]xy]\,\right\}.$$

9.16. Exercise. *Let* $\mathcal{U}$ *be* $\mathbb{R}^m$ *with the metric defined by a norm. Show that* $\mathcal{U}$ *is a complete length* CAT *space if and only if* $\mathcal{U} \overset{iso}{=\!=} \mathbb{E}^m$.

9.17. Exercise. *Assume* $\mathcal{U}$ *is a geodesic* CAT(0) *space. Show that for any two geodesic paths* $\gamma, \sigma\colon [0,1] \to \mathcal{U}$ *the function*

$$t \mapsto |\gamma(t) - \sigma(t)|$$

is convex.

9.18. Proposition. *Assume* $\kappa < \mathrm{K}$. *Then any complete length* $\mathrm{CAT}(\kappa)$ *space is* $\mathrm{CAT}(\mathrm{K})$.

Moreover a space $\mathcal{U}$ *is* $\mathrm{CAT}(\kappa)$ *if* $\mathcal{U}$ *is* $\mathrm{CAT}(\mathrm{K})$ *for all* $\mathrm{K} > \kappa$.

Proof. The first statement follows from Corollary 9.11, the adjacent-angles comparison (9.14a) and the monotonicity of the function $\kappa \mapsto \tilde\measuredangle^\kappa(x\,{}^y_z)$ (1.1d).

The second statement follows since the function $\kappa \mapsto \tilde\measuredangle^\kappa(x\,{}^y_z)$ is continuous. □

D Thin triangles

In this section we define thin triangles and use them to characterize CAT spaces. Inheritance for thin triangles with respect to decomposition is the main result of this section. It will lead to two fundamental constructions: Alexandrov's patchwork globalization (9.29) and Reshetnyak gluing (9.38).

9.19. Definition of κ-thin triangles. *Let* $[x^1x^2x^3]$ *be a triangle of perimeter* $< 2{\cdot}\varpi^\kappa$ *in a metric space. Consider its model triangle* $[\tilde x^1\tilde x^2\tilde x^3] = \tilde\triangle^\kappa(x^1x^2x^3)$ *and the natural map* $[\tilde x^1\tilde x^2\tilde x^3] \to [x^1x^2x^3]$ *that sends a point* $\tilde z \in [\tilde x^i\tilde x^j]$ *to the corresponding point* $z \in [x^ix^j]$ *(that is, such that* $|\tilde x^i - \tilde z| = |x^i - z|$ *and therefore* $|\tilde x^j - \tilde z| = |x^j - z|$*).*

We say the triangle $[x^1x^2x^3]$ *is* κ-*thin if the natural map* $[\tilde x^1\tilde x^2\tilde x^3] \to$ $\to [x^1x^2x^3]$ *is short.*

9.20. Exercise. *Let* $\mathcal{U}$ *be a* ϖ^κ*-geodesic* $\mathrm{CAT}(\kappa)$ *space. Let* $[xyz]$ *be a triangle in* $\mathcal{U}$ *and* $[\tilde x\tilde y\tilde z]$ *be its model triangle in* $\mathbb{M}^2(\kappa)$. *Prove that the natural map* $f\colon [\tilde x\tilde y\tilde z] \to [xyz]$ *is distance-preserving if and only if one of the following conditions hold:*

a) $\measuredangle[x\,{}^{y}_{z}] = \tilde\measuredangle^{\kappa}(x\,{}^{y}_{z})$,
b) $|x-w| = |\tilde x - \tilde w|$ *for some* $\tilde w \in]\tilde y \tilde z[$ *and* $w = f(\tilde w)$,
c) $|v-w| = |\tilde v - \tilde w|$ *for some* $\tilde v \in]\tilde x \tilde y[$, $\tilde w \in]\tilde x \tilde z[$ *and* $v = f(\tilde v)$, $w = f(\tilde w)$.

9.21. Proposition. *Let* $\mathcal{U}$ *be a* ϖ^{κ}*-geodesic space. Then* $\mathcal{U}$ *is* $\mathrm{CAT}(\kappa)$ *if and only if every triangle of perimeter* $< 2\cdot\varpi^{\kappa}$ *in* $\mathcal{U}$ *is* κ*-thin.*

Proof. "If" follows from point-on-side comparison 9.14*b*. "Only if" follows from the angle-sidelength monotonicity 9.15*a*. □

9.22. Corollary. *Suppose* $\mathcal{U}$ *is a* ϖ^{κ}*-geodesic* $\mathrm{CAT}(\kappa)$ *space. Then any local geodesic in* $\mathcal{U}$ *of length* $< \varpi^{\kappa}$ *is length-minimizing.*

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

Since triangle $[\gamma(0)\gamma(a)\gamma(b)]$ is κ-thin, we have

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

for all small $\varepsilon > 0$, a contradiction. □

Now let us formulate the main result of this section. The inheritance lemma states that in any metric space, a triangle is κ-thin if it decomposes into κ-thin triangles. In contrast, $\mathrm{CBB}(\kappa)$ comparisons are not inherited in this way.

9.23. Inheritance lemma. *In a metric space, consider a triangle* $[pxy]$ *that* d e c o m p o s e s *into two triangles* $[pxz]$ *and* $[pyz]$*; that is,* $[pxz]$ *and* $[pyz]$ *have common side* $[pz]$*, and the sides* $[xz]$ *and* $[zy]$ *together form the side* $[xy]$ *of* $[pxy]$*.*

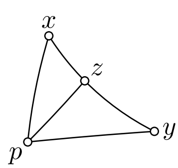

If the triangle $[pxy]$ *has perimeter* $< 2\cdot\varpi^{\kappa}$ *and both triangles* $[pxz]$ *and* $[pyz]$ *are* κ*-thin, then triangle* $[pxy]$ *is* κ*-thin.*

The following model-space lemma is an intermediate statement in the proof of [115, Lemma 2].

9.24. Lemma. *Let* $[\tilde p \tilde x \tilde y]$ *be a triangle in* $\mathbb{M}^2(\kappa)$ *and* $\tilde z \in [\tilde x \tilde y]$*. Consider the solid triangle* $\tilde D = \mathrm{Conv}[\tilde p \tilde x \tilde y]$*. Construct points* $\dot p, \dot x, \dot z, \dot y \in \mathbb{M}^2(\kappa)$ *such that*

$$|\dot p - \dot x| = |\tilde p - \tilde x|, \qquad |\dot p - \dot y| = |\tilde p - \tilde y|, \qquad |\dot p - \dot z| \leqslant |\tilde p - \tilde z|,$$
$$|\dot x - \dot z| = |\tilde x - \tilde z|, \qquad |\dot y - \dot z| = |\tilde y - \tilde z|,$$

where points $\dot{x}$ and $\dot{y}$ lie on either side of $[\dot{p}\dot{z}]$. Set

$$\dot{D} = \mathrm{Conv}[\dot{p}\dot{x}\dot{z}] \cup \mathrm{Conv}[\dot{p}\dot{y}\dot{z}].$$

Then there is a short map $F\colon \tilde{D} \to \dot{D}$ that maps $\tilde{p}$, $\tilde{x}$, $\tilde{y}$ and $\tilde{z}$ to $\dot{p}$, $\dot{x}$, $\dot{y}$ and $\dot{z}$ respectively.

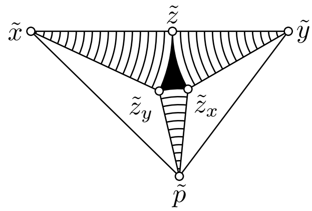

Proof. By Alexandrov's lemma (6.2), there are nonoverlapping triangles $[\tilde{p}\tilde{x}\tilde{z}_y] \stackrel{iso}{=\!=} [\dot{p}\dot{x}\dot{z}]$ and $[\tilde{p}\tilde{y}\tilde{z}_x] \stackrel{iso}{=\!=} [\dot{p}\dot{y}\dot{z}]$ inside triangle $[\tilde{p}\tilde{x}\tilde{y}]$.

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

By construction, F satisfies the conditions of the lemma. □

Proof of 9.23. Construct model triangles $[\dot{p}\dot{x}\dot{z}] = \tilde{\triangle}^{\kappa}(pxz)$ and $[\dot{p}\dot{y}\dot{z}] = \tilde{\triangle}^{\kappa}(pyz)$ so that $\dot{x}$ and $\dot{y}$ lie on opposite sides of $[\dot{p}\dot{z}]$.

Suppose

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

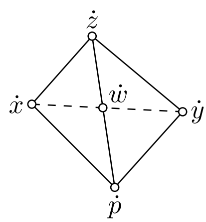

Then for some point $\dot{w} \in [\dot{p}\dot{z}]$, we have

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

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

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

contradicting the triangle inequality.

Thus

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

By Alexandrov's lemma (6.2), this is equivalent to

❶ $$\tilde{\measuredangle}^{\kappa}(x\,{}^{p}_{z}) \leqslant \tilde{\measuredangle}^{\kappa}(x\,{}^{p}_{y}).$$

Let $[\tilde p\tilde x\tilde y]=\tilde\triangle^{\kappa}(pxy)$ and $\tilde z\in[\tilde x\tilde y]$ correspond to z; that is, $|x-{}$ ${}-z|=|\tilde x-\tilde z|$. Inequality ❶ is equivalent to $|p-z|\leqslant|\tilde p-\tilde z|$. Hence Lemma 9.24 applies. Therefore there is a short map F that sends $[\tilde p\tilde x\tilde y]$ to $\dot D=\mathrm{Conv}[\dot p\dot x\dot z]\cup\mathrm{Conv}[\dot p\dot y\dot z]$ in such a way that $\tilde p\mapsto\dot p$, $\tilde x\mapsto\dot x$, $\tilde z\mapsto\dot z$ and $\tilde y\mapsto\dot y$.

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. □

E Function comparison

In this section we give analytic and geometric ways of viewing the point-on-side comparison (9.14*b*) as a convexity condition.

First we obtain a corresponding differential inequality for the distance function in $\mathcal U$. In particular, a geodesic space $\mathcal U$ is CAT(0) if and only if for any $p\in\mathcal U$, the function $\mathrm{dist}_p^2:\mathcal U\to\mathbb R$ is 2-convex; see Section 5F for the definition.

9.25. Theorem. *Suppose $\mathcal U$ is a ϖ^{κ}-geodesic space. Then the following are equivalent:*

a) $\mathcal U$ is $\mathrm{CAT}(\kappa)$,

b) for any $p\in\mathcal U$, the function $f=\mathrm{md}^{\kappa}\circ\mathrm{dist}_p$ satisfies

$$f''+\kappa\cdot f\geqslant 1$$

in $\mathrm B(p,\varpi^{\kappa})$.

Proof. Fix a sufficiently short geodesic $[xy]$ in $\mathrm B(p,\varpi^{\kappa})$. We can assume that the model triangle $[\tilde p\tilde x\tilde y]=\tilde\triangle^{\kappa}(pxy)$ is defined. Let

$$\tilde r(t)=|\tilde p-\mathrm{geod}_{[\tilde x\tilde y]}(t)|,\qquad r(t)=|p-\mathrm{geod}_{[xy]}(t)|.$$

Let $\tilde f=\mathrm{md}^{\kappa}\circ\tilde r$ and $f=\mathrm{md}^{\kappa}\circ r$. By Property 1.1*a*, we have $\tilde f''=1-\kappa\cdot\tilde f$. Clearly $\tilde f(t)$ and $f(t)$ agree at $t=0$ and $t=|x-y|$. The point-on-side comparison (9.14*b*) is the condition $r(t)\leqslant\tilde r(t)$ for all $t\in[0,|x-y|]$. Since md^{κ} is increasing on $[0,\varpi^{\kappa})$, then $r\leqslant\tilde r$ and $f\leqslant\tilde f$ are equivalent. Thus the claim follows by Jensen's inequality (5.14*c*). □

9.26. Corollary. *Suppose $\mathcal U$ is a ϖ^{κ}-geodesic* $\mathrm{CAT}(\kappa)$ *space. Then any ball (closed or open) of radius $R<\frac{\varpi^{\kappa}}{2}$ in $\mathcal U$ is convex.*

Moreover, any open ball of radius $\frac{\varpi^{\kappa}}{2}$ is convex and any closed ball of radius $\frac{\varpi^{\kappa}}{2}$ is ϖ^{κ}-convex.

Proof. Suppose $p \in \mathcal{U}$, $R \leqslant \varpi^\kappa/2$, and two points x and y lie in $\overline{\mathrm{B}}[p,R]$ or $\mathrm{B}(p,R)$. By the triangle inequality, if $|x-y| < \varpi^\kappa$, then any geodesic $[xy]$ lies in $\mathrm{B}(p,\varpi^\kappa)$.

By the function comparison (9.25), the geodesic $[xy]$ lies in $\overline{\mathrm{B}}[p,R]$ or $\mathrm{B}(p,R)$ respectively.

Thus any ball (closed or open) of radius $R < \frac{\varpi^\kappa}{2}$ is ϖ^κ-convex. This implies convexity unless there is a pair of points in the ball at distance at least ϖ^κ. By the triangle inequality, the latter is possible only for the closed ball of radius $\frac{\varpi^\kappa}{2}$. □

Recall that Busemann functions are defined in Proposition 6.1. The following exercise is analogous to Exercise 8.24.

9.27. Exercise. *Let $\mathcal{U}$ be a complete length* $\mathrm{CAT}(\kappa)$ *space and* $\mathrm{bus}_\gamma\colon \mathcal{U} \to \mathbb{R}$ *be the Busemann function for a half-line $\gamma\colon [0,\infty) \to \mathcal{L}$.*

a) If $\kappa = 0$, then the Busemann function bus_γ *is convex.*

b) If $\kappa = -1$, then the function

$$f = \exp \circ \mathrm{bus}_\gamma$$

satisfies

$$f'' - f \geqslant 0.$$

F Development

Geometrically, the development construction (8.25) translates distance comparison into a local convexity statement for subsets of $\mathbb{M}^2(\kappa)$. Recall that a curve in $\mathbb{M}^2(\kappa)$ is (locally) concave with respect to p if (locally) its supergraph with respect to p is a convex subset of $\mathbb{M}^2(\kappa)$; see Definition 8.26.

9.28. Development criterion. *For a ϖ^κ-geodesic space $\mathcal{U}$, the following statements hold:*

a) For any $p \in \mathcal{U}$ and any geodesic $\gamma\colon [0,T] \to \mathrm{B}(p,\varpi^\kappa)$, suppose the κ-development $\tilde\gamma$ in $\mathbb{M}^2(\kappa)$ of γ with respect to p is locally concave. Then $\mathcal{U}$ is $\mathrm{CAT}(\kappa)$.

b) If $\mathcal{U}$ is $\mathrm{CAT}(\kappa)$, *then for any $p \in \mathcal{U}$ and any geodesic $\gamma\colon [0,T] \to \mathcal{U}$ such that the triangle $[p\gamma(0)\gamma(T)]$ has perimeter $< 2\cdot\varpi^\kappa$, the κ-development $\tilde\gamma$ in $\mathbb{M}^2(\kappa)$ of γ with respect to p is concave.*

Proof; (a). Let $\gamma = \mathrm{geod}_{[xy]}$ and $T = |x-y|$. Let $\tilde\gamma\colon [0,T] \to \mathbb{M}^2(\kappa)$ be the concave κ-development based at $\tilde p$ of γ with respect to p. Let us show that the function

❶ $$t \mapsto \measuredangle^\kappa(x\,{}^{p}_{\gamma(t)})$$

is nondecreasing.

For a partition $0 = t^0 < t^1 < \ldots < t^n = T$, let

$$\tilde{y}^i = \tilde{\gamma}(t^i) \quad \text{and} \quad \tau^i = |\tilde{y}^0 - \tilde{y}^1| + |\tilde{y}^1 - \tilde{y}^2| + \ldots + |\tilde{y}^{i-1} - \tilde{y}^i|.$$

Since $\tilde{\gamma}$ is locally concave, for a sufficiently fine partition the broken geodesic $\tilde{y}^0\tilde{y}^1 \ldots \tilde{y}^n$ is locally convex with respect to $\tilde{p}$. Alexandrov's lemma (6.2), applied inductively to pairs of triangles $\tilde{\triangle}^\kappa\{\tau^{i-1}, |p-\tilde{y}^0|, |p-\tilde{y}^{i-1}|\}$ and $\tilde{\triangle}^\kappa\{|\tilde{y}^{i-1} - \tilde{y}^i|, |p - \tilde{y}^{i-1}|, |p - \tilde{y}^i|\}$, shows that the sequence $\tilde{\measuredangle}^\kappa\{|p - \tilde{y}^i|; |p - \tilde{y}^0|, \tau^i\}$ is nondecreasing.

Taking finer partitions and passing to the limit,

$$\max_i\{|\tau^i - t^i|\} \to 0,$$

we get ❶ and the point-on-side comparison (9.14*b*) follows.

(b). Consider a partition $0 = t^0 < t^1 < \ldots < t^n = T$, and let $x^i = \gamma(t^i)$. Construct a chain of model triangles $[\tilde{p}\tilde{x}^{i-1}\tilde{x}^i] = \tilde{\triangle}^\kappa(px^{i-1}x^i)$ with the direction of $[\tilde{p}\tilde{x}^i]$ turning counterclockwise as i grows. By the angle comparison (9.14*c*),

❷ $$\measuredangle[\tilde{x}^i\,{}^{\tilde{x}^{i-1}}_{\tilde{p}}] + \measuredangle[\tilde{x}^i\,{}^{\tilde{x}^{i+1}}_{\tilde{p}}] \geqslant \pi.$$

Since γ is a geodesic,

❸ $$\operatorname{length}\gamma = \sum_{i=1}^{n} |x^{i-1} - x^i| \leqslant |p - x^0| + |p - x^n|.$$

By repeated application of Alexandrov's lemma (6.2), and inequality ❸,

$$\sum_{i=1}^{n} \measuredangle[\tilde{p}\,{}^{\tilde{x}^{i-1}}_{\tilde{x}^i}] \leqslant \tilde{\measuredangle}^\kappa(p\,{}^{x^0}_{x^n}) \leqslant \pi.$$

Then by ❷, the broken geodesic $\tilde{p}\tilde{x}^0\tilde{x}^1 \ldots \tilde{x}^n$ are concave with respect to $\tilde{p}$.

Note that under finer partitions, the broken geodesics $\tilde{x}^0\tilde{x}^1 \ldots \tilde{x}^n$ approach the development of γ with respect to p. Since the broken geodesics are convex, their lengths converge to the length of γ. Hence the result. □

G Patchwork globalization

If $\mathcal{U}$ is a CAT(κ) space, then it is locally CAT(κ). The converse does not hold even for complete length space. For example, $\mathbb{S}^1$ is locally isometric to $\mathbb{R}$, and so is locally CAT(0), but it is easy to find a quadruple of points in $\mathbb{S}^1$ that violates CAT(0) comparison.

The following theorem was essentially proved by Alexandrov [11, Satz 9]; it gives a global condition on geodesics that is necessary and sufficient for a locally CAT(κ) space to be globally CAT(κ). The proof uses thin-triangle decompositions and the inheritance lemma (9.23).

9.29. Patchwork globalization theorem. *For a complete length space $\mathcal{U}$, the following two statements are equivalent:*

a) $\mathcal{U}$ is CAT(κ)*.*

b) $\mathcal{U}$ is locally CAT(κ)*; moreover, pairs of points in $\mathcal{U}$ at distance $< \varpi^\kappa$ are joined by unique geodesics, and these geodesics depend continuously on their endpoint pairs.*

Note that the implication $(a)\Rightarrow(b)$ follows from Theorem 9.8.

9.30. Corollary. *Let $\mathcal{U}$ be a complete length space and $\Omega \subset \mathcal{U}$ be an open locally* CAT(κ) *subset. Then for any point $p \in \Omega$ there is $R > 0$ such that $\overline{\mathrm{B}}[p, R]$ is a convex subset of $\mathcal{U}$ and $\overline{\mathrm{B}}[p, R]$ is* CAT(κ)*.*

Proof. Fix $R > 0$ such that CAT(κ) comparison holds in $\mathrm{B}(p, R)$.

We may assume that $\mathrm{B}(p, R) \subset \Omega$ and $R < \varpi^{\mathrm{K}}$. The same argument as in the proof of the theorem on uniqueness of geodesics (9.8) shows that any two points in $\overline{\mathrm{B}}[p, \frac{R}{2}]$ can be joined by a unique geodesic that depends continuously on the endpoints.

The same argument as in the proof of Corollary 9.26 shows that $\overline{\mathrm{B}}[p, \frac{R}{2}]$ is a convex set. Then $(b)\Rightarrow(a)$ of the patchwork globalization theorem implies that $\overline{\mathrm{B}}[p, \frac{R}{2}]$ is CAT(κ). □

The proof of patchwork globalization uses the following construction:

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

$$(t, s) \mapsto \gamma_t(s)$$

is a line-of-sight map for α with respect to p.

Proof of 9.29. It only remains to prove $(b)\Rightarrow(a)$.

Let $[pxy]$ be a triangle of perimeter $< 2\cdot\varpi^\kappa$ in $\mathcal{U}$. According to Propositions 9.21 and 9.18, it is sufficient to show the triangle $[pxy]$ is κ-thin.

Since pairs of points at distance $< \varpi^\kappa$ are joined by unique geodesics and these geodesics depend continuously on their endpoint pairs, there is a unique and continuous line-of-sight map for $[xy]$ with respect to p.

For a partition

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

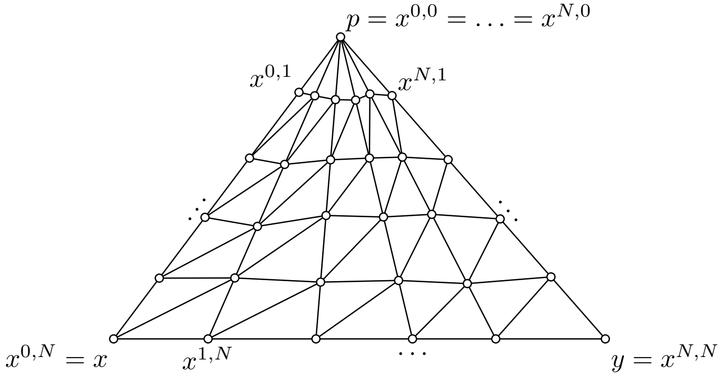

let $x^{i,j} = \gamma_{t^i}(t^j)$. Since the line-of-sight map is continuous, we may assume each triangle $[x^{i,j}x^{i,j+1}x^{i+1,j+1}]$ and $[x^{i,j}x^{i+1,j}x^{i+1,j+1}]$ is κ-thin (see Proposition 9.21).

Now we show that the κ-thin property propagates to $[pxy]$, by repeated application of the inheritance lemma (9.23):

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

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

- ◇ Applying the lemma again shows that the triangles $[xx^{0,N}x^{2,N}]$, $[xx^{0,N}x^{3,N}]$, and so on are κ-thin.

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

The following exercise implies that if the space is proper, then one can drop the condition on continuous dependence of geodesics in the formulation of patchwork globalization.

9.32. Exercise.

a) Suppose pairs of points in a geodesic space $\mathcal{U}$ are joined by unique geodesics. Show that if $\mathcal{U}$ is proper, then these geodesics depend continuously on their endpoint pairs.

b) Construct an example of a complete geodesic space $\mathcal{U}$ such that pairs of points in $\mathcal{U}$ are joined by unique geodesics, but these geodesics do not depend continuously on their endpoint pairs.

H Angles

Recall that ω denotes a selective nonprincipal ultrafilter on $\mathbb{N}$, see Section 3B.

9.33. Angle semicontinuity. *Suppose $\mathcal{U}_n$ is a ϖ^κ-geodesic* CAT(κ) *space for each $n\in\mathbb{N}$ and $\mathcal{U}_n\to\mathcal{U}_\omega$ as $n\to\omega$. Assume that a sequence of hinges $[p_n\,{}^{x_n}_{y_n}]$ in $\mathcal{U}_n$ converges to a hinge $[p_\omega\,{}^{x_\omega}_{y_\omega}]$ in $\mathcal{U}_\omega$. Then*

$$\measuredangle[p_\omega\,{}^{x_\omega}_{y_\omega}]\geqslant\lim_{n\to\omega}\measuredangle[p_n\,{}^{x_n}_{y_n}].$$

Proof. By the angle-sidelength monotonicity (9.15),

$$\measuredangle[p_\omega\,{}^{x_\omega}_{y_\omega}]=\inf\left\{\,\tilde\measuredangle^\kappa(p_\omega\,{}^{\bar x_\omega}_{\bar y_\omega}) : \bar x_\omega\in\,]p_\omega x_\omega],\ \bar y_\omega\in\,]p_\omega y_\omega]\,\right\}.$$

For fixed $\bar x_\omega\in\,]p_\omega x_\omega]$ and $\bar y_\omega\in\,]p_\omega x_\omega]$, choose $\bar x_n\in\,]px_n]$ and $\bar y_n\in$ $\in\,]py_n]$ so that $\bar x_n\to\bar x_\omega$ and $\bar y_n\to\bar y_\omega$ as $n\to\omega$. Clearly

$$\tilde\measuredangle^\kappa(p_n\,{}^{\bar x_n}_{\bar y_n})\to\tilde\measuredangle^\kappa(p_\omega\,{}^{\bar x_\omega}_{\bar y_\omega})$$

as $n\to\omega$.

By the angle comparison (9.14c), $\measuredangle[p_n\,{}^{x_n}_{y_n}]\leqslant\tilde\measuredangle^\kappa(p_n\,{}^{\bar x_n}_{\bar y_n})$. Hence the result. □

Now we verify that the first variation formula holds in the CAT setting. Compare it to the first variation inequality (6.6) which holds for general metric spaces and to the strong angle lemma (8.41) for CBB spaces.

9.34. Strong angle lemma. *Let $\mathcal{U}$ be a ϖ^κ-geodesic* CAT(κ) *space. Then for any hinge $[p\,{}^q_y]$ in $\mathcal{U}$, we have*

❶ $$\measuredangle[p\,{}^q_y]=\lim_{\bar y\to p}\left\{\,\tilde\measuredangle^\kappa(p\,{}^q_{\bar y}) : \bar y\in\,]py]\,\right\}$$

for any $\kappa\in\mathbb{R}$ such that $|p-q|<\varpi^\kappa$.

Proof.

By angle-sidelength monotonicity (9.15), the right-hand side is defined and bigger than or equal to the left-hand side.

By Lemma 6.3, we may take $\kappa=0$ in ❶. By the cosine law and the first variation inequality (6.6), the right-hand side is less than or equal to the left-hand side. □

9.35. First variation. *Let $\mathcal{U}$ be a ϖ^κ-geodesic* CAT(κ) *space. For any nontrivial geodesic $[py]$ in $\mathcal{U}$ and point $q\neq p$ such that $|p-q|<\varpi^\kappa$, we have*

$$|q-\mathrm{geod}_{[py]}(t)|=|q-p|-t\cdot\cos\measuredangle[p\,{}^q_y]+o(t).$$

Proof. The first variation equation is equivalent to the strong angle lemma (9.34), as follows from the Euclidean cosine law. □

9.36. First variation (both-endpoints version). *Let $\mathcal{U}$ be a ϖ^κ-geodesic* CAT(κ) *space. For any nontrivial geodesics $[py]$ and $[qz]$ in $\mathcal{U}$ such that $p \neq q$ and $|p-q| < \varpi^\kappa$, we have*

$$|\mathrm{geod}_{[py]}(t) - \mathrm{geod}_{[qz]}(\tau)| = |q-p| - t\cdot \cos\measuredangle[p\,{}^q_y] - \tau\cdot\cos\measuredangle[q\,{}^p_z] + o(t+\tau).$$

Proof. By the first variation equation (9.35),

$$\begin{aligned}
&|\mathrm{geod}_{[py]}(t) - \mathrm{geod}_{[qz]}(\tau)| =\\
&= |q - \mathrm{geod}_{[py]}(t)| - \tau\cdot\cos\measuredangle[q\,{}^{\mathrm{geod}_{[py]}(t)}_z] + o(\tau)\\
&= |q-p| - t\cdot\cos\measuredangle[p\,{}^q_y] + o(t) - \tau\cdot\cos\measuredangle[q\,{}^{\mathrm{geod}_{[py]}(t)}_z] + o(\tau) =\\
&= |q-p| - t\cdot\cos\measuredangle[p\,{}^q_y] - \tau\cdot\cos\measuredangle[q\,{}^p_z] + o(t+\tau).
\end{aligned}$$

Here the final equality follows from

$$❷\qquad \lim_{t\to 0}\measuredangle[q\,{}^{\mathrm{geod}_{[py]}(t)}_z] = \measuredangle[q\,{}^p_z].$$

The angle semicontinuity (9.33) implies "$\leqslant$" in ❷, and "$\geqslant$" holds by the triangle inequality for angles, since angle comparison (9.14c) gives

$$\lim_{t\to 0}\measuredangle[q\,{}^{p}_{\mathrm{geod}_{[py]}(t)}] = 0.$$

□

We have given elementary proofs of the first-variation statements 9.34, 9.35 and 9.36. Note however that the no-conjugate-point theorem 9.44 not only provides proofs of these statements but also extends the statements from geodesics in CAT(κ) spaces to local geodesics in locally CAT(κ) spaces as follows:

9.37. First variation for local geodesics. *Let $\mathcal{U}$ be a ϖ^κ-geodesic* CAT(κ) *space. For any nontrivial geodesics $[py]$ and $[qz]$ in $\mathcal{U}$ such that $p \neq q$, and any local geodesic $\gamma\colon [0,1] \to \mathcal{U}$ from p to q of length $< \varpi^\kappa$, and any continuous family $\gamma_t\colon [0,1]\to\mathcal{U}$ of local geodesics with $\gamma_0 = \gamma$, $\gamma_t(0) = \mathrm{geod}_{[py]}(t)$ and $\gamma_t(1) = \mathrm{geod}[qz]$, we have*

$$|\mathrm{geod}_{[py]}(t) - \mathrm{geod}_{[qz]}(\tau)| = |q-p| - t\cdot \cos\measuredangle[p\,{}^q_y] - \tau\cdot\cos\measuredangle[q\,{}^p_z] + o(t+\tau).$$

I Reshetnyak gluing theorem

The following theorem was proved by Yuriy Reshetnyak [115], assuming $\mathcal{U}^1$, $\mathcal{U}^2$ are proper and complete. In the following form, the theorem appears in the book of Martin Bridson and André Haefliger [25].

9.38. Reshetnyak gluing theorem. *Suppose $\mathcal{U}^1$ and $\mathcal{U}^2$ are ϖ^κ-geodesic spaces with isometric complete ϖ^κ-convex sets $A_i \subset \mathcal{U}^i$. Let $\iota\colon A_1 \to A_2$ be an isometry. Let $\mathcal{W} = \mathcal{U}^1 \sqcup_\iota \mathcal{U}^2$; that is, $\mathcal{W}$ is the gluing of $\mathcal{U}^1$ and $\mathcal{U}^2$ along ι (see Section 2E).*

Then:

a) Both canonical mappings $\jmath_i\colon \mathcal{U}^i \to \mathcal{W}$ are distance-preserving and the images $\jmath_i(\mathcal{U}^i)$ are ϖ^κ-convex subsets in $\mathcal{W}$.

b) If $\mathcal{U}^1, \mathcal{U}^2$ are CAT(κ) *spaces, then $\mathcal{W}$ is a* CAT(κ) *space.*

Proof. Part (a) follows directly from ϖ^κ-convexity of A_i.

(b). According to (a), we can identify $\mathcal{U}^i$ with its image $\jmath_i(\mathcal{U}^i)$ in $\mathcal{W}$; in this way, the subsets $A_i \subset \mathcal{U}^i$ will be identified and denoted further by A. Thus $A = \mathcal{U}^1 \cap \mathcal{U}^2 \subset \mathcal{W}$, and A is ϖ^κ-convex in $\mathcal{W}$.

Part (b) can be reformulated as follows:

9.39. Reformulation of 9.38*b*. *Let $\mathcal{W}$ be a length space having two ϖ^κ-convex subsets $\mathcal{U}^1, \mathcal{U}^2 \subset \mathcal{W}$ such that $\mathcal{W} = \mathcal{U}^1 \cup \mathcal{U}^2$. Assume the subset $A = \mathcal{U}^1 \cap \mathcal{U}^2$ is complete and ϖ^κ-convex in $\mathcal{W}$, and $\mathcal{U}^1$, $\mathcal{U}^2$ are* CAT(κ) *spaces. Then $\mathcal{W}$ is a* CAT(κ) *space.*

❶ *If $\mathcal{W}$ is ϖ^κ-geodesic, then $\mathcal{W}$ is* CAT(κ).

Indeed, according to 9.21, it is sufficient to show that any triangle $[x^0x^1x^2]$ of perimeter $< 2\cdot\varpi^\kappa$ in $\mathcal{W}$ is κ-thin. This is obviously true if all three points x^0, x^1, x^2 lie in a single $\mathcal{U}^i$. Thus, without loss of generality, we may assume that $x^0 \in \mathcal{U}^1$ and $x^1, x^2 \in \mathcal{U}^2$.

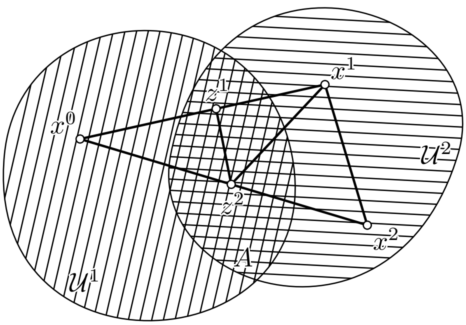

Choose points $z^1, z^2 \in A = \mathcal{U}^1 \cap \mathcal{U}^2$ lying respectively on the sides $[x^0x^1]$, $[x^0x^2]$. Note that all distances between any pair of points from x^0, x^1, x^2, z^1, z^2 are less than ϖ^κ. Therefore

◇ triangle $[x^0z^1z^2]$ lies in $\mathcal{U}^1$,

◇ both triangles $[x^1z^1z^2]$ and $[x^1z^2x^2]$ lie in $\mathcal{U}^2$.

In particular each triangle $[x^0z^1z^2]$, $[x^1z^1z^2]$, $[x^1z^2x^2]$ is κ-thin.

Applying the inheritance lemma for thin triangles (9.23) twice, we get that $[x^0x^1z^2]$ and consequently $[x^0x^1x^2]$ is κ-thin. △

❷ $\mathcal{W}$ *is* CAT(κ) *if* $\kappa \leqslant 0$.

By ❶ it suffices to prove that $\mathcal{W}$ is geodesic.

For $p_1 \in \mathcal{U}^1$, $p_2 \in \mathcal{U}^2$, we may choose a sequence $z^i \in A$ such that $|p_1 - z^i| + |p_2 - z^i|$ converges to $|p_1 - p_2|$, and $|p_1 - z^i|$ and $|p_2 - z^i|$ converge. Since A is complete, it suffices to show z^i is a Cauchy sequence. In that case, the limit point z of z^i satisfies $|p_1 - z| + |p_2 - z| = |p_1 - p_2|$, so the geodesics $[p_1 z]$ in $\mathcal{U}^1$ and $[p_2 z]$ in $\mathcal{U}^2$ together give a geodesic $[p_1 p_2]$ in $\mathcal{U}$.

Suppose z^i is not a Cauchy sequence. Then there are subsequences x^i and y^i of z^i satisfying $\lim |x^i - y^i| > 0$. Let m^i be the midpoint of $[x^i y^i]$. Since $|p_1 - m^i| + |p_2 - m^i| \geqslant |p_1 - p_2|$, and $|p_1 - x^i| + |p_2 - x^i|$ and $|p_1 - y^i| + |p_2 - y^i|$ converge to $|p_1 - p_2|$, then for any $\varepsilon > 0$, we may assume (taking subsequences and possibly relabeling p_1 and p_2)

$$|p_1 - m^i| \geqslant |p_1 - x^i| - \varepsilon, |p_1 - m^i| \geqslant |p_1 - y^i| - \varepsilon.$$

Since triangle $[p^1 x^i y^i]$ is thin, the analogous inequalities hold for the Euclidean model triangle $[\tilde{p}^1 \tilde{x}^i \tilde{y}^i]$. Then there is a nondegenerate limit triangle $[pxy]$ in the Euclidean plane satisfying $|p - x| = |p - y| \leqslant |p - m|$ where m is the midpoint of $[xy]$. This contradiction proves the claim. △

Finally suppose $\kappa > 0$; by rescaling, take $\kappa = 1$. Consider the Euclidean cones $\mathrm{Cone}\,\mathcal{U}^i$ (see Section 6E). By Theorem 11.7*a*, $\mathrm{Cone}\,\mathcal{U}^i$ is a CAT(0) space for $i = 1, 2$.

Geodesics contained in the complement of the tip of $\mathrm{Cone}\,\mathcal{U}^i$ project to geodesics of length $< \pi$ in $\mathcal{U}^i$. It follows that $\mathrm{Cone}\,A$ is convex in $\mathrm{Cone}\,\mathcal{U}^1$ and $\mathrm{Cone}\,\mathcal{U}^2$. By the cone distance formula, $\mathrm{Cone}\,A$ is complete since A is complete.

Gluing along $\mathrm{Cone}\,A$ and applying ❶ and ❷ for $\kappa = 0$, we find that $\mathrm{Cone}\,\mathcal{W}$ is a CAT(0) space. By Theorem 11.7*a*, $\mathcal{W}$ is a CAT(1) space. □

9.40. Exercise. *Let Q be the nonconvex subset of the plane bounded by two half-lines γ_1 and γ_2 with a common starting point and angle α between them. Assume $\mathcal{U}$ is a complete length* CAT(0) *space and γ'_1, γ'_2 are two half-lines in $\mathcal{U}$ with a common starting point and angle α between them. Show that the space glued from Q and $\mathcal{U}$ along the corresponding half-lines is a* CAT(0) *space.*

9.41. Exercise. *Suppose $\mathcal{U}$ is a complete length* CAT(0) *space and $A \subset \subset \mathcal{U}$ is a closed subset. Assume that the doubling of $\mathcal{U}$ in A is* CAT(0). *Show that A is a convex set of $\mathcal{U}$.*

9.42. Exercise. *Let $\mathcal{U}$ be a complete length* $\mathrm{CAT}(1)$ *space and $K \subset \mathcal{U}$ be a closed π-convex set. Assume $K \subset \overline{\mathrm{B}}[p, \frac{\pi}{2}]$ for some point $p \in K$. Show that there is a decreasing continuous one-parameter family of closed convex sets K_t for $t \in [0,1]$ such that $K_0 = \overline{\mathrm{B}}[p, \frac{\pi}{2}]$ and $K_1 = K$.*

(Decreasing means with respect to inclusion; that is $K_{t_0} \supset K_{t_1}$ if $t_0 \leqslant t_1$. Continuous means with respect to Hausdorff distance; that is $K_t \xrightarrow[\mathrm{H}]{} K_{t_0}$ as $t \to t_0$.)

J Space of geodesics

In this section we prove a "no-conjugate-point" theorem for spaces with upper curvature bounds and derive from it a number of statements about local geodesics. These statements will be used in the proof of the Hadamard–Cartan theorem (9.61) and the lifting globalization theorem (9.48), in much the same way as the exponential map is used in Riemannian geometry.

9.43. Proposition. *Let $\mathcal{U}$ be a locally* $\mathrm{CAT}(\kappa)$ *space. Let $\gamma_n \colon [0,1] \to \mathcal{U}$ be 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. Fix $t \in [0,1]$. By Corollary 9.30, we may choose R satisfying $0 < R < \varpi^{\mathrm{K}}$, and such that the ball $\mathcal{B} = \mathrm{B}(\gamma_\infty(t), R)$ is a convex subset of $\mathcal{U}$ and forms a $\mathrm{CAT}(\kappa)$ space.

A local geodesic segment with length less than $R/2$ that intersects $\mathrm{B}(\gamma_\infty(t), R/2)$ cannot leave $\mathcal{B}$, and hence is minimizing by Corollary 9.22. In particular, for all sufficiently large n, if subsegment of γ_n has length less than $R/2$ and contains $\gamma_n(t)$, then it is a geodesic.

Since $\mathcal{B}$ is $\mathrm{CAT}(\kappa)$, geodesic segments in $\mathcal{B}$ depend uniquely and continuously on their endpoint pairs by Theorem 9.8. Thus there is a subinterval $\mathbb{I}$ of $[0,1]$ that contains a neighborhood of t in $[0,1]$ and such that $\gamma_n|_{\mathbb{I}}$ is minimizing for all large n. It follows that the restriction $\gamma_\infty|_{\mathbb{I}}$ is a geodesic, and therefore γ_∞ is a local geodesic. □

The following theorem was proved by the first author and Richard Bishop [3]. In analogy with Riemannian geometry, the main statement of the following theorem could be restated as: In a space of curvature $\leqslant \kappa$, two points cannot be conjugate along a local geodesic of length $< \varpi^\kappa$.

9.44. No-conjugate-point theorem. *Suppose $\mathcal{U}$ is a locally complete, length, locally* $\mathrm{CAT}(\kappa)$ *space. Let $\gamma \colon [0,1] \to \mathcal{U}$ be a local geodesic path* *with length* $< \varpi^\kappa$. *Then for some neighborhoods* $\Omega^0 \ni \gamma(0)$ *and* $\Omega^1 \ni \gamma(1)$, *there is a unique continuous map from the direct product* $\Omega^0 \times \Omega^1 \times [0,1]$ *to* $\mathcal{U}$,

$$(x, y, t) \mapsto \gamma_{xy}(t),$$

such that $\gamma_{xy}\colon [0,1] \to \mathcal{U}$ *is a local geodesic path with* $\gamma_{xy}(0) = x$ *and* $\gamma_{xy}(1) = y$ *for each* $(x,y) \in \Omega^0 \times \Omega^1$, *and the family* γ_{xy} *contains* γ. *Moreover, we can assume that the map*

$$(x, y, t) \mapsto \gamma_{xy}(t)\colon \Omega^0 \times \Omega^1 \times [0,1] \to \mathcal{U}$$

is ℓ*-Lipschitz for any* $\ell > \max\left\{ \frac{\operatorname{sn}^\kappa r}{\operatorname{sn}^\kappa \ell} : 0 \leqslant r \leqslant \ell \right\}$.

The following lemma was suggested to us by Alexander Lytchak. The proof proceeds by piecing together CAT(κ) neighborhoods of points on a curve to construct a new CAT(κ) space. Exercise 9.68 is inspired by the original idea of the proof of the no-conjugate-point theorem (9.44) given in [3].

9.45. Patchwork along a curve. *Let* $\mathcal{U}$ *be a locally complete, length, locally* CAT(κ) *space, and* $\alpha\colon [a,b] \to \mathcal{U}$ *be a curve.*

Then there is a complete length CAT(κ) *space* $\mathcal{N}$ *with an open set* $\hat{\Omega} \subset \mathcal{N}$, *a curve* $\hat{\alpha}\colon [0,1] \to \hat{\Omega}$, *and an open locally isometric immersion* $\Phi\colon \hat{\Omega} \looparrowright \mathcal{U}$ *such that* $\Phi \circ \hat{\alpha} = \alpha$.

Moreover if α *is simple, then one can assume in addition that* Φ *is an open embedding; thus* $\hat{\Omega}$ *is locally isometric to a neighborhood of* $\Omega = \Phi(\hat{\Omega})$ *of* α.

Proof. According to Corollary 9.30, for any $t \in [a,b]$, there is $r(t) > 0$ such that the closed ball $\overline{\mathrm{B}}[\alpha(t), r(t)]$ is a convex set that forms a complete length CAT(κ) space.

Choose balls $\mathcal{B}^i = \overline{\mathrm{B}}[\alpha(t^i), r(t^i)]$ for some partition $a = t^0 < t^1 < \dots$ $\dots < t^n = b$ in such a way that $\operatorname{Int} \mathcal{B}^i \supset \alpha([t^{i-1}, t^i])$ for all $i > 0$.

Consider the disjoint union $\bigsqcup_i \mathcal{B}^i = \left\{ (i, x) : x \in \mathcal{B}^i \right\}$ with the minimal equivalence relation $\sim$ such that $(i, x) \sim (i-1, x)$ for all $i > 0$. Let $\mathcal{N}$ be the space obtained by gluing the $\mathcal{B}^i$ along $\sim$. Note that $A^i = \mathcal{B}^i \cap$ $\cap\, \mathcal{B}^{i-1}$ is convex in $\mathcal{B}^i$ and in $\mathcal{B}^{i-1}$. Applying the Reshetnyak gluing theorem (9.38) n times, we conclude that $\mathcal{N}$ is a complete length CAT(κ) space.

For $t \in [t^{i-1}, t^i]$, let $\hat{\alpha}(t)$ be the equivalence class of $(i, \alpha(t))$ in $\mathcal{N}$. Let $\hat{\Omega}$ be the ε-neighborhood of $\hat{\alpha}$ in $\mathcal{N}$, where $\varepsilon > 0$ is chosen so that $\mathrm{B}(\alpha(t), \varepsilon) \subset \mathcal{B}^i$ 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 $\Phi\colon \mathcal{N} \to \mathcal{U}$, $\hat{\alpha}\colon [0,1] \to \mathcal{N}$ and $\hat{\Omega} \subset \mathcal{N}$ satisfy the conclusion of the main part of the lemma.

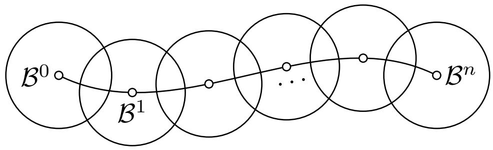

To prove the final statement in the lemma, we only have to choose $\varepsilon > 0$ so that in addition, $|\alpha(\tau) - \alpha(\tau')| > 2\cdot\varepsilon$ if $\tau \leqslant t^{i-1}$ and $t^i \leqslant \tau'$ for some i. □

Proof of 9.44. Apply patchwork along γ (9.45). □

The No-conjugate-point theorem (9.44) allows us to move a local geodesic so that its endpoints follow given trajectories. The following corollary describes how this process might terminate.

9.46. Corollary. *Let $\mathcal{U}$ be a locally complete, length, locally* CAT(κ) *space. Suppose $\gamma\colon [0,1] \to \mathcal{U}$ is a local geodesic with length $< \varpi^\kappa$. Let $\alpha^i\colon [0,1] \to \mathcal{U}$, for $i = 0, 1$, be curves starting at $\gamma(0)$ and $\gamma(1)$ respectively.*

Then there is a uniquely determined pair consisting of an interval $\mathbb{I}$ satisfying $0 \in \mathbb{I} \subset [0,1]$, and a continuous family of local geodesics $\gamma_t\colon [0,1] \to \mathcal{U}$ for $t \in \mathbb{I}$, such that

a) $\gamma_0 = \gamma$, $\gamma_t(0) = \alpha^0(t)$, $\gamma_t(1) = \alpha^1(t)$, and γ_t has length $< \varpi^\kappa$,

b) if $\mathbb{I} \neq [0,1]$, then $\mathbb{I} = [0,a)$, where either γ_t converges uniformly to a local geodesic γ_a of length ϖ^κ, or for some fixed $s \in [0,1]$ the curve $\gamma_t(s) : [0,a) \to \mathcal{U}$ is a Lipschitz curve with no limit as $t \to a-$.

Proof. Uniqueness follows from Theorem 9.44.

Let $\mathbb{I}$ be the maximal interval for which there is a family γ_t satisfying condition (a). By Theorem 9.44, such an interval exists and is open in $[0,1]$. Suppose $\mathbb{I} \neq [0,1]$. Then $\mathbb{I} = [0,a)$ for some $0 < a \leqslant 1$. It suffices to show that $\mathbb{I}$ satisfies condition (b).

For each fixed $s \in [0,1]$, define the curve $\alpha_s : [0,a) \to \mathcal{U}$ by $\alpha_s(t) = \gamma_t(s)$. By Theorem 9.44, α_s is ℓ-Lipschitz for some ℓ.

If α_s for some value of s does not converge as $t \to a-$, then condition (b) holds. If each α_s converges as $t \to a-$, then γ_t converges as $t \to a-$, say to γ_a. By Proposition 9.43, γ_a is a local geodesic and

$$\operatorname{length}\gamma_t \to \operatorname{length}\gamma_a \leqslant \varpi^\kappa.$$

By maximality of $\mathbb{I}$, $\operatorname{length}\gamma_a = \varpi^\kappa$ and so condition (b) again holds. □

9.47. Corollary. *Let $\mathcal{U}$ be a complete locally* CAT(κ) *length space, and $\alpha\colon [0,1] \to \mathcal{U}$ be a path of length $< \varpi^\kappa$ that starts at p and ends at q. Then:*

a) *There is a unique homotopy of local geodesic paths* $\gamma_t\colon [0,1]\to \to \mathcal{U}$ *such that* $\gamma_0(t)=\gamma_t(0)=p$ *and* $\gamma_t(1)=\alpha(t)$ *for any* t.

b) *For any* $t\in[0,1]$,

$$\operatorname{length}\gamma_t\leqslant \operatorname{length}(\alpha|_{[0,t]}),$$

and equality holds for given t *if and only if the restriction* $\alpha|_{[0,t]}$ *is a reparametrization of* γ_t.

Moreover, instead of completeness of $\mathcal{U}$, *one can assume that*

$$W=\{\,x\in\mathcal{U}\;:\;|x-p|+|x-q|\leqslant\ell\,\}$$

is complete.

Proof. By Corollary 9.46, substituting $\bar q = p$ for q and taking $\alpha^0(t)=p$ and $\alpha^1(t)=\alpha(t)$ for all $t\in[0,1]$, there is an interval $\mathbb{I}$ such that (a) holds for all $t\in\mathbb{I}$, and either $\mathbb{I}=[0,1]$ or $\mathbb{I}=[0,a)$ for some $a\leqslant 1$.

By patchwork along a curve (9.45), the values of t for which condition (b) holds form an open subset of $\mathbb{I}$ containing 0; clearly this subset is also closed in $\mathbb{I}$. Therefore (b) holds on all of $\mathbb{I}$.

Corollary 9.46 implies that $\mathbb{I}=[0,1]$. Indeed if $\mathbb{I}=[0,a)$, then either $\operatorname{length}\gamma_t\to\varpi^\kappa$ as $t\to a-$, or for some fixed $s\in[0,1]$ the Lipschitz curve $\gamma_t(s)\,:\,[0,a)\to\mathcal{U}$ has no limit as $t\to a-$. Since $\operatorname{length}\alpha<\varpi^\kappa$, 9.46 implies that neither of these is possible. □

K Lifting globalization

The Hadamard–Cartan theorem (9.61) states that the universal metric cover of a complete locally CAT(0) space is CAT(0). The lifting globalization theorem gives an appropriate generalization of the above statement to arbitrary curvature bounds; it could be also described as a global version of Gauss's lemma.

9.48. Lifting globalization theorem. *Suppose* $\mathcal{U}$ *is a complete length locally* CAT(κ) *space and* $p\in\mathcal{U}$. *Then there is a complete* CAT(κ) *length space* $\mathcal{B}$, *with a point* $\hat p$ *such that there is a locally isometric map* $\Phi\colon\mathcal{B}\to\to\mathcal{U}$ *such that* $\Phi(\hat p)=p$ *and the following lifting property holds: for any path* $\alpha\colon[0,1]\to\mathcal{U}$ *with* $\alpha(0)=p$ *and* $\operatorname{length}\alpha<\varpi^\kappa/2$, *there is a unique path* $\hat\alpha\colon[0,1]\to\mathcal{B}$ *such that* $\hat\alpha(0)=\hat p$ *and* $\Phi\circ\hat\alpha\equiv\alpha$.

Note that the lifting property implies that $\Phi(\mathcal{B})\supset \mathrm{B}(p,\varpi^\kappa/2)$ and by completeness $\Phi(\mathcal{B})\supset\overline{\mathrm{B}}[p,\varpi^\kappa/2]$. Also since $\mathcal{B}$ is CAT(κ), the closed ball $\overline{\mathrm{B}}[\hat p,\tfrac{\varpi^\kappa}{2}]_{\mathcal{B}}$ is a weakly convex set in $\mathcal{B}$ (see 9.26); in particular $\overline{\mathrm{B}}[\hat p,\tfrac{\varpi^\kappa}{2}]_{\mathcal{B}}$

is a complete length CAT(κ) space. Therefore we can assume in addition that $|\hat p-\hat x|\leqslant\varpi^\kappa/2$ for any $\hat x\in\mathcal{B}$; or equivalently

$$\overline{\mathrm{B}}[\hat p,\tfrac{\varpi^\kappa}{2}]_{\mathcal{B}}=\mathcal{B}.$$

Before proving the theorem we state and prove its corollary.

9.49. Corollary. *Suppose $\mathcal{U}$ is a complete length locally* CAT(κ) *space. Then for any $p\in\mathcal{U}$ there is $\rho_p>0$ such that $\overline{\mathrm{B}}[p,\rho_p]$ is a complete length* CAT(κ) *space.*

Moreover, we can assume that $\rho_p<\frac{\varpi^\kappa}{2}$ for any p and the function $p\mapsto\rho_p$ is 1-Lipschitz.

Proof. Assume $\Phi\colon\mathcal{B}\to\mathcal{U}$ and $\hat p\in\mathcal{B}$ are provided by the lifting globalization theorem (9.48).

Since Φ is local isometry, we can choose $r>0$ so that the restriction of Φ to $\overline{\mathrm{B}}[\hat p,r]$ is distance-preserving. By the lifting globalization, the image $\Phi(\overline{\mathrm{B}}[\hat p,r])$ coincides with the ball $\overline{\mathrm{B}}[p,r]$. This proves the first part of the theorem.

To prove the second part, let us choose ρ_p to be the maximal value $\leqslant\frac{\varpi^\kappa}{2}$ such that $\overline{\mathrm{B}}[p,\rho_p]$ is a complete length CAT(κ) space. By Corollary 9.26, the ball

$$\overline{\mathrm{B}}[q,\rho_p-|p-q|]$$

is convex in $\overline{\mathrm{B}}[p,\rho_p]$. Therefore

$$\overline{\mathrm{B}}[q,\rho_p-|p-q|]$$

is a complete length CAT(κ) space for any $q\in\mathrm{B}(p,\rho_p)$. In particular, $\rho_q\geqslant\rho_p-|p-q|$ for any $p,q\in\mathcal{U}$. Hence the second statement follows. ☐

The proof of the lifting globalization theorem relies heavily on the properties of the space of local geodesic paths discussed in Section 9J. The following lemma proved by the first author and Richard Bishop [2] is a key step in the proof.

9.50. Radial lemma. *Let $\mathcal{U}$ be a length locally* CAT(κ) *space, and suppose $p\in\mathcal{U}$, $R\leqslant\varpi^\kappa$. Assume the ball $\overline{\mathrm{B}}[p,\bar R]$ is complete for any $\bar R<R$, and there is a unique geodesic path, $\mathrm{path}_{[px]}$, from p to any point $x\in\mathrm{B}(p,R)$ that depends continuously on x. Then $\mathrm{B}(p,\frac{R}{2})$ is a ϖ^κ-geodesic* CAT(κ) *space.*

Proof. Without loss of generality, we may assume $\mathcal{U}=\mathrm{B}(p,R)$.

Set $f=\mathrm{md}^\kappa\circ\mathrm{dist}_p$. Let us show that

❶ $$f''+\kappa\cdot f\geqslant 1.$$

Fix $z \in \mathcal{U}$. We will apply Theorem 9.44 for the unique geodesic path γ from p to z. The notations Ω^0, Ω^1, γ_{xy}, $\mathcal{N}$, $\hat{x}$, $\hat{y}$ will be as in Theorem 9.48; in particular, $z \in \Omega^1$.

By assumption, $\gamma_{py} = \mathrm{path}_{[py]}$ for any $y \in \Omega^1$. Consequently, $f(y) = \mathrm{md}^\kappa|\hat{p} - \hat{y}|_{\mathcal{N}}$. Applying the function comparison (9.25) in $\mathcal{N}$, we have that $f'' + \kappa \cdot f \geqslant 1$ in Ω^1; whence ❶ follows. △

Fix $r < \frac{R}{2}$. Proving the following claim takes most of the remaining proof:

❷ $\overline{\mathrm{B}}[p, r]$ *is a convex set in* $\mathcal{U}$.

Choose arbitrary $x, z \in \overline{\mathrm{B}}[p, r]$. First note that ❶ implies the following claim.

❸ *If* $\gamma\colon [0, 1] \to \mathcal{U}$ *is a local geodesic path from* x *to* z *and* $\mathrm{length}\,\gamma < \varpi^\kappa$, *then* $\mathrm{length}\,\gamma \leqslant 2\cdot r$ *and* γ *lies completely in* $\overline{\mathrm{B}}[p, r]$.

Note that $|x - z| < \varpi^\kappa$. Thus to prove Claim ❷, it is sufficient to show that there is a geodesic path from x to z. Note that by assumption $\overline{\mathrm{B}}[p, 2\cdot r]$ is complete. Therefore Corollary 9.47 implies the following:

❹ *Given a path* $\alpha\colon [0, 1] \to \mathcal{U}$ *from* x *to* z *with* $\mathrm{length}\,\alpha < 2\cdot r$, *there is a local geodesic path* γ *from* x *to* z *such that*

$$\mathrm{length}\,\gamma \leqslant \mathrm{length}\,\alpha.$$

Further, let us prove the following:

❺ *There is a unique local geodesic path* γ_{xz} *in* $\overline{\mathrm{B}}[p, r]$ *from* x *to* z.

Denote by Δ_{xz} the set of all local geodesic paths in $\overline{\mathrm{B}}[p, r]$ from x to z. By Corollary 9.46, there is a bijection $\Delta_{xz} \to \Delta_{pp}$. According to ❶, Δ_{pp} contains only the constant path. Claim ❺ follows.

Note that claims ❸, ❹ and ❺ imply that γ_{xz} is minimizing; hence Claim ❷.

Further, Claim ❸ and the no-conjugate-point theorem (9.44) together imply that the map $(x, z) \mapsto \gamma_{xz}$ is continuous.

Therefore by the patchwork globalization theorem (9.29), $\overline{\mathrm{B}}[p, r]$ is a ϖ^κ-geodesic CAT(κ) space.

Since

$$\mathrm{B}(p, R) = \bigcup_{r<R} \overline{\mathrm{B}}[p, r],$$

then $\mathrm{B}(p, R)$ is convex in $\mathcal{U}$ and CAT(κ) comparison holds for any quadruple in $\mathrm{B}(p, R)$. Therefore $\mathrm{B}(p, \varpi^\kappa/2)$ is CAT(κ). □

In the following proof, we construct a space $\mathfrak{G}_p$ of local geodesic paths that start at p. The space $\mathfrak{G}_p$ comes with a marked point $\hat{p}$ and the endpoint map $\Phi\colon \mathfrak{G}_p \to \mathcal{U}$. One can think of the map Φ as an analog of $\exp_p$ in the Riemannian case; in this case the space $\mathfrak{G}_p$ is the ball of radius ϖ^κ in the tangent space at p, equipped with the metric pulled back by $\exp_p$.

We are going to set $\mathcal{B} = \mathrm{B}(\hat{p}, \varpi^\kappa/2) \subset \mathfrak{G}_p$, and use the radial lemma (9.50) to prove that $\mathcal{B}$ is a ϖ^κ-geodesic CAT(κ) space.

Proof of 9.48. Suppose $\hat{\gamma}$ is a homotopy of local geodesic paths that start at p. Thus the map

$$\hat{\gamma}\colon (t,\tau) \mapsto \hat{\gamma}_t(\tau)\colon [0,1]\times[0,1] \to \mathcal{U}$$

is continuous, and the following holds for each t:

- ◇ $\hat{\gamma}_t(0) = p$,
- ◇ $\hat{\gamma}_t\colon [0,1] \to \mathcal{U}$ is a local geodesic path in $\mathcal{U}$.

Denote by $\vartheta(\hat{\gamma})$ the length traced by the ends of $\hat{\gamma}_t$; that is, $\vartheta(\hat{\gamma})$ is the length of the path $t \mapsto \hat{\gamma}_t(1)$.

Let $\mathfrak{G}_p$ be the set of all local geodesic paths with length $< \varpi^\kappa$ in $\mathcal{U}$ that start at p. Denote by $\hat{p} \in \mathfrak{G}_p$ the constant path $\hat{p}(t) \equiv p$. Given $\alpha, \beta \in \mathfrak{G}_p$, define

$$|\alpha - \beta|_{\mathfrak{G}_p} = \inf_{\hat{\gamma}}\{\vartheta(\hat{\gamma})\},$$

with the exact lower bound taken along all homotopies $\hat{\gamma}\colon [0,1]\times[0,1] \to \to \mathcal{U}$ such that $\hat{\gamma}_0 = \alpha$, $\hat{\gamma}_1 = \beta$ and $\hat{\gamma}_t \in \mathfrak{G}_p$ for all $t \in [0,1]$.

From 9.44, we have $|\alpha - \beta|_{\mathfrak{G}_p} > 0$ for distinct α and β; that is,

❻ $|{*} - {*}|_{\mathfrak{G}_p}$ *is a metric on* $\mathfrak{G}_p$*.*

Further, again from Theorem 9.44, we have

❼ *The map*

$$\Phi\colon \xi \mapsto \xi(1)\colon \mathfrak{G}_p \to \mathcal{U}$$

is a local isometry. In particular, $\mathfrak{G}_p$ *is locally* CAT(κ)*.*

Let $\alpha\colon [0,1] \to \mathcal{U}$ be a path with $\operatorname{length}\alpha < \varpi^\kappa$ and $\alpha(0) = p$. The homotopy constructed in Corollary 9.47 can be regarded as a path in $\mathfrak{G}_p$, say $\hat{\alpha}\colon [0,1] \to \mathfrak{G}_p$, such that $\hat{\alpha}(0) = \hat{p}$ and $\Phi \circ \hat{\alpha} = \alpha$; in particular $\hat{\alpha}_t(1) \equiv \alpha(t)$ for any t. By ❼,

$$\operatorname{length}(\hat{\alpha})_{\mathfrak{G}_p} = \operatorname{length}(\alpha)_{\mathcal{U}}.$$

Moreover, it follows that α is a local geodesic path of $\mathcal{U}$ if and only if $\hat{\alpha}$ is a local geodesic path of $\mathfrak{G}_p$.

Further, from Corollary 9.47, for any $\xi \in \mathfrak{G}_p$ and path $\hat\alpha\colon [0,1] \to \mathfrak{G}_p$ from $\hat p$ to ξ, we have

$$\begin{aligned}\operatorname{length}\hat\alpha &= \operatorname{length}\Phi\circ\hat\alpha \geqslant\\ &\geqslant \operatorname{length}\xi =\\ &= \operatorname{length}\hat\xi\end{aligned}$$

where equality holds only if $\hat\alpha$ is a reparametrization of $\hat\xi$. In particular,

❽ $$|\hat p - \xi|_{\mathfrak{G}_p} = \operatorname{length}\xi$$

and $\hat\xi\colon [0,1] \to \mathfrak{G}_p$ is the unique geodesic path from $\hat p$ to ξ. Clearly, the map $\xi \mapsto \hat\xi$ is continuous.

By ❽ and Proposition 9.43,

❾ *For any $\bar R < \varpi^\kappa$, the closed ball $\overline{\mathrm{B}}[\hat p, \bar R]$ in $\mathfrak{G}_p$ is complete.*

Take $\mathrm{B}(\hat p, \varpi^\kappa/2)$ and Φ constructed above. According to the radial lemma (9.50), $\mathrm{B}(\hat p, \varpi^\kappa/2)$ is a ϖ^κ-geodesic CAT(κ) space. The map Φ extends to its completion $\mathcal{B} = \overline{\mathrm{B}}[\hat p, \varpi^\kappa/2]$. All the remaining statements are already proved. □

L Reshetnyak majorization

9.51. Definition. *Let $\mathcal{X}$ be a metric space, $\tilde\alpha$ be a simple closed curve of finite length in $\mathbb{M}^2(\kappa)$, and $D \subset \mathbb{M}^2(\kappa)$ be a closed region bounded by $\tilde\alpha$. A length-nonincreasing map $F\colon D \to \mathcal{X}$ is called majorizing if it is length-preserving on $\tilde\alpha$.*

In this case, we say that D majorizes the curve $\alpha = F\circ\tilde\alpha$ under the map F.

The following proposition is a consequence of the definition.

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

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

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 $\tilde{s} = \tilde{r}$. □

9.53. Corollary. *Let $[pxy]$ be a triangle of perimeter $< 2\cdot\varpi^\kappa$ in a metric space $\mathcal{X}$. Assume a convex region $D \subset \mathbb{M}^2(\kappa)$ majorizes $[pxy]$. Then $D = \mathrm{Conv}[\tilde{p}\tilde{x}\tilde{y}]$ for a model triangle $[\tilde{p}\tilde{x}\tilde{y}] = \tilde{\triangle}^\kappa(pxy)$, and 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.

9.54. Majorization theorem. *Any closed curve α with length smaller than $2\cdot\varpi^\kappa$ in a ϖ^κ-geodesic $\mathrm{CAT}(\kappa)$ space is majorized by a convex region in $\mathbb{M}^2(\kappa)$.*

This theorem was proved by Yuriy Reshetnyak [115]; our proof uses a trick that we learned from the lectures of Werner Ballmann [15]. Another proof can be built on Kirszbraun's theorem (10.14), but it works only for complete spaces.

The case when α is a triangle, say $[pxy]$, is the base and is nontrivial. In this case, by Proposition 9.53, the majorizing convex region has to be isometric to $\mathrm{Conv}[\tilde{p}\tilde{x}\tilde{y}]$, where $[\tilde{p}\tilde{x}\tilde{y}] = \tilde{\triangle}^\kappa(pxy)$. There is a majorizing map for $[pxy]$ whose image W is the image of the line-of-sight map (definition 9.31) for $[xy]$ from p, but 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 nonconvex at z (as in the picture). Equip $\mathcal{Q}$ with the length metric. Then $\mathcal{Q}$ is CAT(0) by Reshetnyak gluing (9.38). For triangle $[pxy]_{\mathcal{Q}}$ in $\mathcal{Q}$ and its model triangle $[\tilde{p}\tilde{x}\tilde{y}]$ in $\mathbb{E}^2$, we have

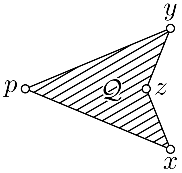

$$|\tilde{x} - \tilde{y}| = |x - y|_{\mathcal{Q}} = |x - z| + |z - y|.$$

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, by the first variation formula (8.41), for $\varepsilon = 1 - s$ we have

$$|\tilde{x} - \tilde{w}| = |\tilde{x} - \tilde{\gamma}_{\frac{1}{2}}(s)| = |x - z| + o(\varepsilon)$$

and

$$|x - F(\tilde{w})| = |x - \gamma_{\frac{1}{2}}(s)| = |x - z| - \varepsilon\cdot\cos\measuredangle[z\,{}^{p}_{x}] + o(\varepsilon).$$

Thus F is not majorizing.

In the following proofs, $x^1 \dots x^n$ ($n \geqslant 3$) denotes a broken geodesic with vertexes $x^1, \dots, x^n$, 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.

Our first lemma gives a model space construction based on repeated application of Lemma 9.24. Recall that convex and concave curves with respect to a point are defined in 8.26.

9.55. Lemma. *In $\mathbb{M}^2(\kappa)$, 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. Assume*

$$\operatorname{length}\beta + |p-x| + |p-y| < 2\cdot\varpi^\kappa.$$

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 triangle $[pxy]_D$. 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 broken geodesic β; the general case then follows by approximation. Namely, since β is concave it can be approximated by broken geodesics 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 broken geodesic with $x^0 = x$ and $x^n = y$. Consider a sequence of broken geodesics $\beta_i = x^0x^1 \dots x^{i-1}y_i$ such that $|p - y_i| = |p - y|$ and β_i has same length as β; that is,

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

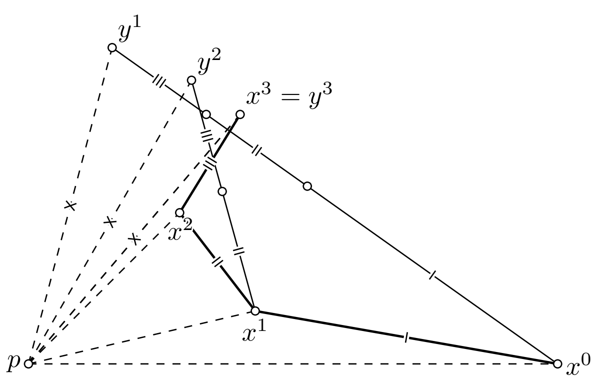

Clearly $\beta_n = \beta$. Sequentially applying Alexandrov's lemma (6.2) shows that each of the broken geodesics $\beta_{n-1}, \beta_{n-2}, \ldots, \beta_1$ is concave with respect to p. Let D_i be the subgraph of β_i with respect to p. Applying Lemma 9.24 gives a short map $G_i\colon D_i \to D_{i+1}$ 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

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

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

9.56. Lemma. *Let $[pxy]$ be a triangle of perimeter $< 2\cdot\varpi^\kappa$ in a ϖ^κ-geodesic* $\mathrm{CAT}(\kappa)$ *space $\mathcal{U}$. In $\mathbb{M}^2(\kappa)$, 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 $\gamma = \mathrm{geod}_{[xy]}$ and $T = |x-y|$. As in the proof of the development criterion (9.28), take a partition

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

and set $x^i = \gamma(t^i)$. Construct a chain of model triangles $[\tilde{p}\tilde{x}^{i-1}\tilde{x}^i] = \tilde{\triangle}^\kappa(px^{i-1}x^i)$, 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 broken geodesic $\tilde{x}^0 \ldots \tilde{x}^n$.

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 $x \in D_n$. Then x lies on or inside some triangle $[\tilde{p}\tilde{x}^{i-1}\tilde{x}^i]$. Define $H_n(x)$ by first mapping x 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 κ-thin (9.21), 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 $y_n \in D_n$ with $y_n \to y \in D$ and $z_n \in D_n$ with $z_n \to z \in D$. Since $\delta_n \to 0$ under increasingly finer partitions and geodesics in $\mathcal{U}$ vary continuously with their endpoints (9.29), we have $H_n(x_n) \to H(x)$ and $H_n(y_n) \to H(y)$. Since

$$|H_n(x_n) - H_n(y_n)| \leqslant |x_n - y_n|_{D_n} + 2{\cdot}\delta_n,$$

where the left-hand side converges to $|H(x) - H(y)|$ and the right-hand side converges to $|x - y|_D$, it follows that H is short. □

Proof of 9.54. We begin by proving the theorem in case α is polygonal.

First suppose α is a triangle, say $[pxy]$. By assumption, the perimeter of $[pxy]$ is less than $2{\cdot}\varpi^\kappa$. This is the base case for the induction.

Let $\tilde\gamma$ be the κ-development of $[xy]$ with respect to p, where $\tilde\gamma$ has basepoint $\tilde p$ and subgraph D. By the development criterion (9.28), $\tilde\gamma$ is concave. By Lemma 9.55, there is a short map $G\colon \operatorname{Conv}\tilde\triangle^\kappa(pxy) \to$ $\to D$. Further, by Lemma 9.56, D majorizes $[pxy]$ under a majorizing map $H\colon D \to \mathcal{U}$. Clearly $H \circ G$ is a majorizing map for $[pxy]$.

Now we claim that any closed n-gon $[x^1x^2\dots x^n]$ of perimeter less than $2{\cdot}\varpi^\kappa$ in a CAT(κ) 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.

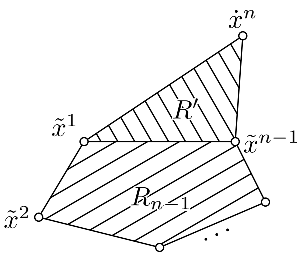

Assume the statement is true for $(n-1)$-gons, $n \geqslant 4$. Then $[x^1x^2\dots x^{n-1}]$ is majorized by a convex polygonal region

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

in $\mathbb{M}^2(\kappa)$ under a map F_{n-1} satisfying $F_{n-1}(\tilde x^i) = x^i$ for all i. Take $\dot x^n \in \mathbb{M}^2(\kappa)$ such that $[\tilde x^1\tilde x^{n-1}\dot x^n] = \tilde\triangle^\kappa(x^1x^{n-1}x^n)$ 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 $[x^1x^{n-1}x^n]$ 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 $[x^1x^2\dots x^n]$ (as in Definition 9.51).

If R is a convex subset of $\mathbb{M}^2(\kappa)$, 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 $>\pi$. 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 (9.38), R carrying its length metric is a CAT(κ)-space. Moreover $[\tilde{x}^{n-2}\tilde{x}^{n-1}]\cup[\tilde{x}^{n-1}\dot{x}^n]$ is a geodesic of R. Thus $[\tilde{x}^1\tilde{x}^2\dots$ $\dots\tilde{x}^{n-1}\dot{x}^n]_R$ is a closed $(n-1)$-gon in R, to which the induction hypothesis applies. The resulting short map from a convex region in $\mathbb{M}^2(\kappa)$ to R, followed by F, is the desired majorizing map.

Note that in fact we have proved the following:

❶ *Let F_{n-1} be a majorizing map for the polygon $[x^1x^2\dots x^{n-1}]$, and $\dot{F}$ be a majorizing map for the triangle $[x^1x^{n-1}x^n]$. Then there is a majorizing map F_n for the polygon $[x^1x^2\dots x^n]$ such that*

$$\operatorname{Im}F_{n+1}=\operatorname{Im}F_n\cup\operatorname{Im}\dot{F}.$$

We now use this claim to prove the theorem for general curves.

Assume $\alpha\colon[0,\ell]\to\mathcal{U}$ is an arbitrary closed curve with natural parameter. Choose a sequence of partitions $0=t_n^0<t_n^2<\dots<t_n^n=\ell$ so that:

- ◇ The set $\{t^i_{n+1}\}_{i=0}^{n+1}$ is obtained from the set $\{t^i_n\}_{i=0}^n$ by adding one element.
- ◇ For some sequence $\varepsilon_n\to0+$, we have $t^i_n-t^{i-1}_n<\varepsilon_n$ for all i.

Inscribe in α a sequence of polygons P_n with vertexes $\alpha(t^i_n)$. Apply the claim above, to get a sequence of majorizing maps $F_n\colon R_n\to\mathcal{U}$. Note that for all $m>n$

- ◇ $\operatorname{Im}F_m$ lies in an ε_n-neighborhood of $\operatorname{Im}F_n$
- ◇ $\operatorname{Im}F_m\backslash\operatorname{Im}F_n$ lies in an ε_n-neighborhood of α.

It follows that the set

$$K=\alpha\cup\left(\bigcup_n\operatorname{Im}F_n\right)$$

is compact. Therefore the sequence (F_n) has a partial limit as $n\to\infty$; say F. Clearly F is a majorizing map for α. □

The following exercise is the rigidity case of the majorization theorem.

9.57. Exercise. *Let $\mathcal{U}$ be a ϖ^κ-geodesic* CAT(κ) *space and $\alpha\colon[0,\ell]\to\mathcal{U}$ be a closed curve with arclength parametrization. Assume that $\ell<2\cdot\varpi^\kappa$ and there is a closed convex curve $\tilde{\alpha}\colon[0,\ell]\to\mathbb{M}^2(\kappa)$ such that*

$$|\alpha(t_0)-\alpha(t_1)|_{\mathcal{U}}=|\tilde{\alpha}(t_0)-\tilde{\alpha}(t_1)|_{\mathbb{M}^2(\kappa)}$$

for any t_0 *and* t_1*. Then there is a distance-preserving map* $F\colon \mathrm{Conv}\,\tilde\alpha\to\to\mathcal{U}$ *such that* $F\colon\tilde\alpha(t)\mapsto\alpha(t)$ *for any* t*.*

9.58. Exercise. *Two majorizations* $F\colon D\to\mathcal{U}$ *and* $F'\colon D'\to\mathcal{U}$ *will be called* equivalent *if* $F'=F\circ\iota$ *for an isometry* $\iota\colon D\to D'$*.*

Show that a closed rectifiable curve in a $\mathrm{CAT}(0)$ *space has an isometric majorization map if and only if the majorization map is unique up to equivalence.*

For $n=1$, the following lemma states that in a $\mathrm{CAT}(\kappa)$ space, a sharp triangle comparison implies the presence of an isometric copy of the convex hull of the model triangle. This rigidity statement was proved by Alexandrov [11].

9.59. Arm lemma. *Let* $\mathcal{U}$ *be a* ϖ^κ*-geodesic* $\mathrm{CAT}(\kappa)$ *space, and* $P=[x^0x^1\dots x^{n+1}]$ *be a polygon of length* $<2\cdot\varpi^\kappa$ *in* $\mathcal{U}$*. Suppose* $\tilde P=[\tilde x^0\tilde x^1\dots\dots\tilde x^{n+1}]$ *is a convex polygon in* $\mathbb{M}^2(\kappa)$ *such that*

$$❷\qquad |\tilde x^i-\tilde x^{i-1}|_{\mathbb{M}^2(\kappa)}=|x^i-x^{i-1}|_{\mathcal{U}}\quad\text{and}\quad \measuredangle[x^i\,{}^{x^{i-1}}_{x^{i+1}}]\geqslant\measuredangle[\tilde x^i\,{}^{\tilde x^{i-1}}_{\tilde x^{i+1}}]$$

for all i*. Then*

a) $|\tilde x^0-\tilde x^{n+1}|_{\mathbb{M}^2(\kappa)}\leqslant|x^0-x^{n+1}|_{\mathcal{U}}$.

b) *Equality holds in (a) if and only if the map* $\tilde x^i\mapsto x^i$ *can be extended to a distance-preserving map of* $\mathrm{Conv}(\tilde x^0,\tilde x^1\dots\tilde x^{n+1})$ *onto* $\mathrm{Conv}(x^0,x^1\dots x^{n+1})$*.*

Proof; a). By majorization (9.54), P is majorized by a convex region $\tilde D$ in $\mathbb{M}^2(\kappa)$. By Proposition 9.52 and the definition of angle, $\tilde D$ is bounded by a convex polygon $\tilde P_R=[\tilde y^0\tilde y^1\dots\tilde y^{n+1}]$ that satisfies

$$|\tilde y^i-\tilde y^{i\pm1}|_{\mathbb{M}^2(\kappa)}=|x^i-x^{i\pm1}|_{\mathcal{U}},\qquad |\tilde y^0-\tilde y^{n+1}|_{\mathbb{M}^2(\kappa)}=|x^0-x^{n+1}|_{\mathcal{U}},$$
$$\measuredangle[\tilde y^i\,{}^{\tilde y^{i-1}}_{\tilde y^{i+1}}]\geqslant\measuredangle[x^i\,{}^{x^{i-1}}_{x^{i+1}}]\geqslant\measuredangle[\tilde x^i\,{}^{\tilde x^{i-1}}_{\tilde x^{i+1}}]$$

for $1\leqslant i\leqslant n$; the last inequality follows from ❷.

The classical arm lemma [118] gives $|\tilde x^0-\tilde x^{n+1}|\leqslant|\tilde y^0-\tilde y^{n+1}|$. Since $|\tilde y^0-\tilde y^{n+1}|=|x^0-x^{n+1}|$, part (a) follows.

b). Suppose equality holds in (a). Then angles at the j-th vertex of $\tilde P$, P, and $\tilde P_R$ are equal for $1\leqslant j\leqslant n$, and we may take $\tilde P=\tilde P_R$.

Let $F\colon\tilde D\to\mathcal{U}$ be the majorizing map for P, where $\tilde D$ is the convex region bounded by $\tilde P$, and $F|_{\tilde P}$ is length-preserving.

❸ *Let* $\tilde x,\tilde y,\tilde z$ *be three vertexes of* $\tilde P$*, and* x,y,z *be the corresponding vertexes of* P*. If* $|\tilde x-\tilde y|=|x-y|$*,* $|\tilde y-\tilde z|=|x-z|$ *and* $\measuredangle[\tilde y\,{}^{\tilde x}_{\tilde z}]=\measuredangle[y\,{}^{x}_{z}]$*, then* $F|_{\mathrm{Conv}(\tilde x,\tilde y,\tilde z)}$ *is distance-preserving.*

Indeed, since F is majorizing, F restricts to distance-preserving maps from $[\tilde x\tilde y]$ to $[xy]$ and $[\tilde y\tilde z]$ to $[yz]$. Suppose $\tilde p\in[\tilde x\tilde y]$ and $\tilde q\in[\tilde y\tilde z]$. Then

❹ $$|\tilde p-\tilde q|_{\mathbb{M}^2(\kappa)}=|F(\tilde p)-F(\tilde q)|_{\mathcal{U}}.$$

This inequality holds in one direction by majorization, and in the other direction by the angle comparison (9.14*c*). By the first variation formula (9.36), it follows that each pair of corresponding angles of triangles $[\tilde x\tilde y\tilde z]$ and $[xyz]$ are equal. But then ❹ holds for p,q on any two sides of these triangles, so F is distance-preserving on every geodesic of $\mathrm{Conv}(\tilde p,\tilde x,\tilde y)$. Hence the claim. △

❺ *Suppose* $F|_{\mathrm{Conv}(\tilde x^{n+1},\tilde x^0,\tilde x^1,\dots,\tilde x^k)}$ *is distance-preserving for some* k, $1\leqslant$ $\leqslant k\leqslant n-1$. *Then* $F|_{\mathrm{Conv}(\tilde x^{n+1},\tilde x^0,\tilde x^1,\dots,\tilde x^{k+1})}$ *is distance-preserving.*

To verify this claim, set

$$\begin{aligned}\tilde p&=[\tilde x^{k-1}\tilde x^{k+1}]\cap[\tilde x^k\tilde x^{n+1}],\\ p&=[x^{k-1}x^{k+1}]\cap[x^kx^{n+1}].\end{aligned}$$

The following maps are distance-preserving:

(i) $F|_{\mathrm{Conv}(\tilde x^{k-1},\tilde x^k,\tilde x^{k+1})}$,

(ii) $F|_{\mathrm{Conv}(\tilde x^{k+1},\tilde x^{k-1},\tilde x^{n+1})}$,

(iii) $F|_{\mathrm{Conv}(\tilde x^{n+1},\tilde x^k,\tilde x^{k+1})}$.

Indeed, (i) follows from ❸. Therefore $|\tilde x^{k-1}-\tilde x^{k+1}|=|x^{k-1}-x^{k+1}|$, and so F restricts to a distance-preserving map from $[\tilde x^{k-1}\tilde x^{k+1}]$ onto $[x^{k-1}x^{k+1}]$. With the induction hypothesis ❺, it follows that $F(\tilde p)=p$, hence

❻ $$\measuredangle[\tilde x^{k-1}\,{}^{\tilde x^{k+1}}_{\tilde x^{n+1}}]=\measuredangle[x^{n+1}\,{}^{x^{k-1}}_{x^{k+1}}].$$

Then (ii) follows from ❻ and ❸. Since $|\tilde x^k-\tilde x^{n+1}|=|x^k-x^{n+1}|$, (iii) follows from ❻ and (i).

Let $\tilde\gamma$ be a geodesic of $\mathrm{Conv}(\tilde x^{n+1},\tilde x^0,\tilde x^1\dots\tilde x^{k+1})$. Then $\mathrm{length}\,\tilde\gamma<$ $<\varpi^\kappa$. If $\tilde\gamma$ does not contain the point $[\tilde x^{k-1}\tilde x^{k+1}]\cap[\tilde x^k\tilde x^{n+1}]$, it follows from the induction hypothesis and (i) (ii), (iii) that $\gamma=F\circ\tilde\gamma$ is a local geodesic of length $<\varpi^\kappa$. By 9.22, γ is a geodesic.

By continuity, $F\circ\tilde\gamma$ is a geodesic for all $\tilde\gamma$; so ❺ follows. The base of the induction is Exercise 9.57. It finishes the proof of part *(b)*. ☐

9.60. Exercise. *Let* $\mathcal{U}$ *be a complete length* CAT(0) *space and suppose for 4 points* $x^1,x^2,x^3,x^4\in\mathcal{U}$ *there is a convex quadrangle* $[\tilde x^1\tilde x^2\tilde x^3\tilde x^4]$ *in* $\mathbb{E}^2$ *such that*

$$|x^i-x^j|_{\mathcal{U}}=|\tilde x^i-\tilde x^j|_{\mathbb{E}^2}$$

for all i *and* j. *Show that* $\mathcal{U}$ *contains an isometric copy of the solid quadrangle* $[\tilde x^1\tilde x^2\tilde x^3\tilde x^4]$; *that is, the convex hull of* $\tilde x^1,\tilde x^2,\tilde x^3,\tilde x^4$ *in* $\mathbb{E}^2$.

M Hadamard–Cartan theorem

The development of Alexandrov geometry was greatly influenced by the Hadamard–Cartan theorem. Its original formulation states that if M is a complete Riemannian manifold with nonpositive sectional curvature, 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, the theorem appeared in the lectures of Élie Cartan [37]. For surfaces in the Euclidean plane, the theorem was proved by Hans von Mangoldt [91], and a few years later independently by Jacques Hadamard [64].

Formulations for metric spaces of different generality were proved by Herbert Busemann [35], Willi Rinow [116], and Michael Gromov [61, p. 119]. A detailed proof of Gromov's statement when $\mathcal{U}$ is proper was given by Werner Ballmann [14], using Birkhoff's curve-shortening. A proof in the non-proper geodesic case was given by the first author and Richard Bishop [3]. This proof applies more generally, to convex spaces (see Exercise 9.68). It was pointed out by Bruce Kleiner [15] and independently by Martin Bridson and André Haefliger [25] that this proof extends to length spaces, as well as geodesic spaces, giving the following statement:

9.61. Hadamard–Cartan theorem. *Let $\mathcal{U}$ be a complete, simply connected length locally* CAT(0) *space. Then $\mathcal{U}$ is* CAT(0).

Proof. Since $\varpi^\kappa = \infty$, the map $\Phi\colon \mathcal{B} \to \mathcal{U}$ in Theorem 9.48 is a metric covering; that is, Φ is a length-preserving covering map. Since $\mathcal{U}$ is simply connected, $\Phi\colon \mathcal{B} \to \mathcal{U}$ is an isometry. □

To formulate the generalized Hadamard–Cartan theorem, we need the following definition.

9.62. Definition. *Given $\ell \in (0, \infty]$, a metric space $\mathcal{X}$ is called ℓ-simply connected if it is connected and any closed curve of length $< \ell$ is null-homotopic in the class of curves of length $< \ell$ in $\mathcal{X}$.*

Note that there is a subtle difference between simply connected and ∞-simply connected spaces; the first states that any closed curve is null-homotopic while the second means that any rectifiable curve is null-homotopic in the class of rectifiable curves. However, as follows from Proposition 9.65, for locally CAT(κ) spaces these two definitions are equivalent. This fact makes it possible to deduce the Hadamard–Cartan theorem directly from the generalized Hadamard–Cartan theorem.

9.63. Generalized Hadamard–Cartan theorem. *A complete length space $\mathcal{U}$ is* $\mathrm{CAT}(\kappa)$ *if and only if $\mathcal{U}$ is $2\cdot\varpi^\kappa$-simply connected and $\mathcal{U}$ is locally* $\mathrm{CAT}(\kappa)$.

For proper spaces, the generalized Hadamard–Cartan theorem was proved by Brian Bowditch [23]. In the proof we need the following lemma.

9.64. Lemma.
Assume $\mathcal{U}$ is a complete length locally $\mathrm{CAT}(\kappa)$ *space, $\varepsilon > 0$, and $\gamma_1, \gamma_2 \colon \mathbb{S}^1 \to \mathcal{U}$ are two closed curves. Assume*

a) $\operatorname{length}\gamma_1, \operatorname{length}\gamma_2 < 2\cdot\varpi^\kappa - 4\cdot\varepsilon$;
b) $|\gamma_1(x) - \gamma_2(x)| < \varepsilon$ *for any $x \in \mathbb{S}^1$, and the geodesic $[\gamma_1(x)\gamma_2(x)]$ is uniquely defined and depends continuously on x;*
c) γ_1 *is majorized by a convex region in* $\mathbb{M}^2(\kappa)$.

Then γ_2 is majorized by a convex region in $\mathbb{M}^2(\kappa)$.

Proof. Let D be a convex region in $\mathbb{M}^2(\kappa)$ that majorizes γ_1 under the map $F \colon D \to \mathcal{U}$ (see Definition 9.51). Denote by $\tilde\gamma_1$ the curve bounding D such that $F \circ \tilde\gamma_1 = \gamma_1$. Since

$$\begin{aligned}\operatorname{length}\tilde\gamma_1 = \operatorname{length}\gamma_1 <{}&\\ &< 2\cdot\varpi^\kappa - 4\cdot\varepsilon,\end{aligned}$$

there is a point $\tilde p \in D$ such that $|\tilde p - \tilde\gamma(x)|_{\mathbb{M}^2(\kappa)} < \frac{\varpi^\kappa}{2} - \varepsilon$ for any $x \in \mathbb{S}^1$. Denote by α_x the join of the paths $F \circ \operatorname{path}_{[p\tilde\gamma_1(x)]_{\mathbb{M}^2(\kappa)}}$ and $\operatorname{path}_{[\gamma_1(x)\gamma_2(x)]}$ in $\mathcal{U}$. Note that α_x depends continuously on x, and

$$\operatorname{length}\alpha_x < \tfrac{\varpi^\kappa}{2} \quad \text{and} \quad \alpha_x(1) = \gamma_2(x)$$

hold for any x.

Let us apply the lifting globalization theorem (9.48) for $p = F(\tilde p)$. We obtain a ϖ^κ-geodesic $\mathrm{CAT}(\kappa)$ space $\mathcal{B}$ and a locally isometric map $\Phi \colon \mathcal{B} \to \mathcal{U}$ with $\Phi(\hat p) = p$ for some $\hat p \in \mathcal{B}$, and with the lifting property for the curves starting at p with length $< \varpi^\kappa/2$. Applying the lifting property for α_x, we get existence of a curve $\hat\gamma_2 \colon \mathbb{S}^1 \to \mathcal{B}$ such that

$$\gamma_2 = \Phi \circ \hat\gamma_2.$$

Since $\mathcal{B}$ is a geodesic $\mathrm{CAT}(\kappa)$ space, we can apply the majorization theorem (9.54) for $\hat\gamma_2$. The composition of the obtained majorization with Φ is a majorization of γ_2. □

Proof of Theorem 9.63. The "only if" part follows from the Reshetnyak majorization theorem (9.54).

Let γ_t, $t \in [0,1]$ be a null-homotopy of curves in $\mathcal{U}$; that is, $\gamma_0(x) = p$ for some $p \in \mathcal{U}$ and any $x \in \mathbb{S}^1$. Assume further that length $\gamma_t < 2\cdot\varpi^\kappa$ for any t. To prove the "if" part, it is sufficient to show that γ_1 is majorized by a convex region in $\mathbb{M}^2(\kappa)$ if $\mathcal{U}$ is locally CAT(κ).

By semicontinuity of length (2.6), we can choose $\varepsilon > 0$ sufficiently small that

$$\operatorname{length}\gamma_t < 2\cdot\varpi^\kappa - 4\cdot\varepsilon$$

for all t.

By Corollary 9.49, we may assume in addition that $\mathrm{B}(\gamma_t(x), \varepsilon)$ is CAT(κ) for any t and x.

Choose a partition $0 = t_0 < t_1 < \ldots < t_n = 1$ so that $|\gamma_{t_i}(x) - \gamma_{t_{i-1}}(x)| < \varepsilon$ for any i and x. According to 9.8, for any i, the geodesic $[\gamma_{t_i}(x)\gamma_{t_{i-1}}(x)]$ depends continuously on x.

Note that $\gamma_0 = \gamma_{t_0}$ is majorized by a convex region in $\mathbb{M}^2(\kappa)$. Applying the lemma n times, we see that the same holds for $\gamma_1 = \gamma_{t_n}$. □

9.65. Proposition. *Let $\mathcal{U}$ be a complete length locally* CAT(κ) *space. Then $\mathcal{U}$ is simply connected if and only if it is ∞-simply connected.*

Proof; "if" part. It is sufficient to show that any closed curve in $\mathcal{U}$ is homotopic to a closed broken geodesic.

Let γ_0 be a closed curve in $\mathcal{U}$. According to Corollary 9.49, there is $\varepsilon > 0$ such that $\mathrm{B}(\gamma(x), \varepsilon)$ is CAT(κ) for any x.

Choose a broken geodesic γ_1 such that $|\gamma_0(x) - \gamma_1(x)| < \varepsilon$ for any x. By 9.8, $\mathrm{path}_{[\gamma_0(x)\gamma_1(x)]}$ is uniquely defined and depends continuously on x.

Hence $\gamma_t(x) = \mathrm{path}_{[\gamma_0(x)\gamma_1(x)]}(t)$ gives a homotopy from γ_0 to γ_1.

"Only if" part. The proof is similar.

Assume γ_t is a homotopy between two rectifiable curves γ_0 and γ_1. Fix $\varepsilon > 0$ so that the ball $\mathrm{B}(\gamma_t(x), \varepsilon)$ is CAT(κ) for any t and x. Choose a partition $0 = t_0 < t_1 < \ldots < t_n = 1$ so that

$$|\gamma_{t_{i-1}}(x) - \gamma_{t_i}(x)| < \tfrac{\varepsilon}{10}$$

for any i and x. Set $\hat\gamma_{t_0} = \gamma_0$, $\hat\gamma_{t_n} = \gamma_{t_n}$. For each $0 < i < n$, approximate γ_{t_i} by a closed broken geodesic $\hat\gamma_i$.

Construct the homotopy from $\hat\gamma_{t_{i-1}}$ to $\hat\gamma_{t_i}$ setting

$$\hat\gamma_t = \mathrm{path}_{[\hat\gamma_{t_{i-1}}(x)\hat\gamma_{t_i}(x)]}(t).$$

Since ε is sufficiently small, by 9.13, we get that

$$\operatorname{length}\hat\gamma_t < 10\cdot(\operatorname{length}\hat\gamma_{t_{i-1}} + \operatorname{length}\hat\gamma_{t_i}).$$

In particular, $\hat{\gamma}_t$ is rectifiable for all t.

Joining the obtained homotopies for all i we obtain a homotopy from γ_0 to γ_1 in the class of rectifiable curves. □

9.66. Exercise. *Let $\mathcal{X}$ be a double cover of $\mathbb{E}^3$ that branches along two distinct lines ℓ and m. Show that $\mathcal{X}$ is* CAT(0) *if and only if ℓ intersects m at a right angle.*

9.67. Exercise. *Let $\mathcal{U}$ be a complete length* CAT(0) *space. Assume $\tilde{\mathcal{U}} \to \mathcal{U}$ is a metric covering branching along a geodesic. Show that $\tilde{\mathcal{U}}$ is* CAT(0).

More generally, assume $A \subset \mathcal{U}$ is a closed convex subset and $f\colon \mathcal{X} \to \to \mathcal{U}\backslash A$ is a metric cover. Denote by $\bar{\mathcal{X}}$ the completion of $\mathcal{X}$, and $\bar{f}\colon \bar{\mathcal{X}} \to \to \mathcal{U}$ the continuous extension of f. Let $\tilde{\mathcal{U}}$ be the space glued from $\bar{\mathcal{X}}$ and A by identifying x and $\bar{f}(x)$ if $\bar{f}(x) \in A$. Show that $\tilde{\mathcal{U}}$ is CAT(0).

About convex spaces. A convex space $\mathcal{X}$ is a geodesic space such that the function $t \mapsto |\gamma(t) - \sigma(t)|$ is convex for any two geodesic paths $\gamma, \sigma : [0,1] \to \mathcal{X}$. A locally convex space is a length space in which every point has a neighborhood that is a convex space in the restricted metric.

9.68. Exercise. *Assume $\mathcal{X}$ is a convex space such that the angle of any hinge is defined. Show that $\mathcal{X}$ is* CAT(0).

The following exercise gives an analog of Hadamard–Cartan theorem for locally convex spaces; see also [3].

9.69. Exercise. *Show that a complete, simply connected, locally convex space is a convex space.*

N Convex sets

Recall that according to Corollary 9.26, any ball (closed or open) of radius $R < \frac{\varpi^\kappa}{2}$ in a ϖ^κ-geodesic CAT(κ) space is convex.

9.70. Proposition. *Any weakly ϖ^κ-convex set in a complete length* CAT(κ) *space is ϖ^κ-convex.*

Proof. Follows from the uniqueness of geodesics in CAT(κ) spaces (9.8). □

9.71. Closest-point projection lemma. *Let $\mathcal{U}$ be a complete length* CAT(κ) *space, and $K \subset \mathcal{U}$ be a closed ϖ^κ-convex set. Assume that $\operatorname{dist}_K p < \frac{\varpi^\kappa}{2}$ for some point $p \in \mathcal{U}$. Then there is a unique point $p^* \in K$ that minimizes the distance to p; that is, $|p^* - p| = \operatorname{dist}_K p$.*

Proof. Fix r properly between $\operatorname{dist}_K p$ and $\frac{\varpi^\kappa}{2}$. By the function comparison (9.25), the function $f = \operatorname{md}^\kappa \circ \operatorname{dist}_p$ is strongly convex in $\overline{\mathrm{B}}[p,r]$.

The lemma follows from Lemma 14.4 applied to the subspace $K' = K \cap \overline{\mathrm{B}}[p,r]$ and the restriction $f|_{K'}$. □

9.72. Exercise. *Let $\mathcal{U}$ be a complete length* CAT(0) *space and $K \subset \mathcal{U}$ be a closed convex set. Show that the closest-point projection $\mathcal{U} \to K$ is short.*

9.73. Advanced exercise. *Let $\mathcal{U}$ be a complete length* CAT(1) *space and $K \subset \mathcal{U}$ be a closed π-convex set. Assume $K \subset \overline{\mathrm{B}}[p, \frac{\pi}{2}]$ for some point $p \in K$. Show that there is a short retraction of $\mathcal{U}$ to K.*

9.74. Proposition. *Let $\mathcal{U}$ be a ϖ^κ-geodesic* CAT(κ) *space and $K \subset \mathcal{U}$ be a closed ϖ^κ-convex set. Let*

$$f = \operatorname{sn}^\kappa \circ \operatorname{dist}_K .$$

Then

$$f'' + \kappa{\cdot}f \geqslant 0$$

holds in $\mathrm{B}(K, \frac{\varpi^\kappa}{2})$.

Proof. It is sufficient to show that Jensen's inequality (5.14*c*) holds on a sufficiently short geodesic $[pq]$ in $\mathrm{B}(K, \frac{\varpi^\kappa}{2})$. Since

$$\operatorname{dist}_K p, \operatorname{dist}_K q < \tfrac{\varpi^\kappa}{2},$$

we may assume that

❶ $$|p-q| + \operatorname{dist}_K p + \operatorname{dist}_K q < \varpi^\kappa.$$

For each $x \in [pq]$, we need to find a value $h(x) \in \mathbb{R}$ such that

$$h(x) \leqslant f(x), \qquad h(p) = f(p), \qquad h(q) = f(q),$$

and

❷ $$h'' + \kappa{\cdot}h \geqslant 0$$

along $[pq]$.

Denote by p^* and q^* the closest-point projections of p and q on K; they are provided by lemma 9.71. From ❶ and the triangle inequality, we have

$$|p^* - q^*| < \varpi^\kappa.$$

Since K is ϖ^κ-convex, $K\supset[pq]$; in particular

$$\operatorname{dist}_K x\leqslant\operatorname{dist}_{[p^*q^*]}x$$

for any $x\in\mathcal{U}$.

There is a majorizing map $F:D\to\mathcal{U}$ for quadrangle $[pp^*q^*q]$, as in Definition 9.51 and the Reshetnyak majorization theorem 9.54. By Proposition 9.52, the figure D is a solid convex model quadrangle $[\tilde p\tilde p^*\tilde q^*\tilde q]$ in $\mathbb{M}^2(\kappa)$ such that

$$\begin{aligned}|\tilde p-\tilde p^*|_{\mathbb{M}^2(\kappa)}&=|p-p^*|_{\mathcal{U}} & |\tilde p-\tilde q|_{\mathbb{M}^2(\kappa)}&=|p-q|_{\mathcal{U}}\\ |\tilde q-\tilde q^*|_{\mathbb{M}^2(\kappa)}&=|q-q^*|_{\mathcal{U}} & |\tilde p^*-\tilde q^*|_{\mathbb{M}^2(\kappa)}&=|p^*-q^*|_{\mathcal{U}}.\end{aligned}$$

Given $x\in[pq]$, denote by $\tilde x$ the corresponding point on $[\tilde p\tilde q]$. Then

$$|[pq]-x|_{\mathcal{U}}\leqslant|[\tilde p\tilde q]-\tilde x|_{\mathbb{M}^2(\kappa)}.$$

Set

$$h(x)=\operatorname{sn}^\kappa|[\tilde p\tilde q]-\tilde x|_{\mathbb{M}^2(\kappa)}.$$

By straightforward calculations, ❷ holds and hence the statement follows. □

9.75. Corollary. *Let $\mathcal{U}$ be a complete length* CAT(κ) *space and $K\subset\mathcal{U}$ be a closed locally convex set. Then there is an open set $\Omega\supset K$ such that the function*

$$f=\operatorname{sn}^\kappa\circ\operatorname{dist}_K$$

satisfies

$$f''+\kappa\cdot f\geqslant 0$$

in Ω.

Proof. Fix $p\in K$. By Corollary 9.26, $\overline{\mathrm{B}}[p,r]$ is convex for all small $r>0$.

Since K is locally convex, there is $r_p>0$ such that the intersection $K'=K\cap\mathrm{B}(p,r_p)$ is convex.

Note that

$$\operatorname{dist}_K x=\operatorname{dist}_{K'}x$$

for any $x\in\mathrm{B}(p,\tfrac{r_p}{2})$. Therefore the statement holds for

$$\Omega=\bigcup_{p\in K}\mathrm{B}(p,\tfrac{r_p}{2}).$$

□

9.76. Theorem. *Assume $\mathcal{U}$ is a complete length* CAT(κ) *space and $K\subset\subset\mathcal{U}$ is a closed connected locally convex set. Assume $|x-y|<\varpi^\kappa$ for any $x,y\in K$. Then K is convex.*

In particular, if $\kappa \leqslant 0$*, then any closed connected locally convex set in a* $\mathcal{U}$ *is convex.*

The following proof is inspired by the answer of Sergei Ivanov to the question of Nathan Reading [71].

Proof. Since K is locally convex, it is locally path-connected. Since K is connected and locally path connected it is path-connected.

Fix two points $x, y \in K$. Let us connect x to y by a path $\alpha\colon [0,1] \to K$. Since $|x-\alpha(s)| < \varpi^\kappa$ for any s, Theorem 9.8 implies that the geodesic $[x\alpha(s)]$ is uniquely defined and depends continuously on s.

Let $\Omega \supset K$ be the open set provided by Corollary 9.75. 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. By Corollary 9.75, the function $f = \mathrm{sn}^\kappa \circ \mathrm{dist}_K \mathcal{U}$ satisfies the differential inequality

❸ $$f'' + \kappa{\cdot}f \geqslant 0$$

along $[x\alpha(s)]$.

Since

$$|x-\alpha(s)| < \varpi^\kappa, \qquad f(x) = f(\alpha(s)) = 0,$$

then the barrier inequality (5.14*b*) implies that $f(z) \leqslant 0$ for $z \in [x\alpha(s)]$; that is $[x\alpha(s)] \subset K$, a contradiction. □

O Remarks

The following question was known in folklore in the 80's, but it seems that in print it was first mentioned Michael Gromov [59, 6.B$_1$(f)]. We do not see any reason why it should be true, but we also cannot construct a counterexample.

9.77. Open question. *Let* $\mathcal{U}$ *be a complete length* CAT(0) *space and* $K \subset \mathcal{U}$ *be a compact set. Is it true that* K *lies in a convex compact set* $\bar{K} \subset \mathcal{U}$*?*

The question can easily be reduced to the case when K is finite; so far it is not even known if any three points in a complete length CAT(0) space lie in a compact convex set.

One of the most beautiful applications of the Reshetnyak gluing theorem is given by Dmitri Burago, Serge Ferleger, and Alexey Kononenko [28–31]. They use it to study billiards; a short survey on the subject is written by Dmitri Burago [26].

Chapter 10

Kirszbraun revisited

This chapter is based on our paper [6] and an earlier paper of Urs Lang and Viktor Schroeder [79].

A Short map extension definitions

10.1. Theorem. *A complete length space $\mathcal{L}$ is* $\mathrm{CBB}(\kappa)$ *if and only if for any 3-point set V_3 and any 4-point set $V_4 \supset V_3$ in $\mathcal{L}$, any short map $f\colon V_3 \to \mathbb{M}^2(\kappa)$ can be extended to a short map $F\colon V_4 \to \mathbb{M}^2(\kappa)$ (so $f = F|_{V_3}$).*

The proof of the "only if" part of Theorem 10.1 can be obtained as a corollary of Kirszbraun's theorem (10.14). We present another, more elementary proof; using the following analog of Alexandrov's lemma (6.2).

We say that two triangles with a common vertex *do not overlap* if their convex hulls intersect only at the common vertex.

10.2. Overlap lemma. *Let $[\tilde{x}^1\tilde{x}^2\tilde{x}^3]$ be a triangle in $\mathbb{M}^2(\kappa)$. Let $\tilde{p}^1, \tilde{p}^2, \tilde{p}^3$ be points in $\mathbb{M}^2(\kappa)$ such that, for any permutation $\{i, j, k\}$ of $\{1, 2, 3\}$, we have*

(i) $|\tilde{p}^i - \tilde{x}^k| = |\tilde{p}^j - \tilde{x}^k|$,

(ii) $\tilde{p}^i$ and $\tilde{x}^i$ lie in the same closed halfspace determined by $[\tilde{x}^j\tilde{x}^k]$,

If no pair of triangles $[\tilde{p}^i\tilde{x}^j\tilde{x}^k]$ overlap, then

$$\measuredangle\tilde{p}^1 + \measuredangle\tilde{p}^2 + \measuredangle\tilde{p}^3 > 2\cdot\pi,$$

where $\measuredangle\tilde{p}^i$ denotes $\measuredangle[\tilde{p}^i\,{}^{\tilde{x}^k}_{\tilde{x}^j}]$ for a permutation $\{i, j, k\}$ of $\{1, 2, 3\}$.

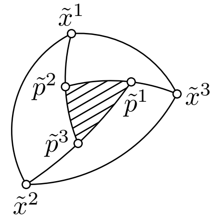

Remarks. If $\kappa \leqslant 0$, then the overlap lemma can be proved without using condition (i). This follows since the sum of external angles for the hexagon $[\tilde p^1\tilde x^2\tilde p^3\tilde x^1\tilde p^2\tilde x^3]$ and its area is $2\cdot\pi - \kappa\cdot a$, where a denotes the area of the hexagon.

The diagram shows that condition (i) is essential in case $\kappa > 0$.

Proof. Rotate the triangle $[\tilde p^3\tilde x^1\tilde x^2]$ around $\tilde x^1$ to make $[\tilde x^1\tilde p^3]$ coincide with $[\tilde x^1\tilde p^2]$. Let $\dot x^2$ denote the image of $\tilde x^2$ after rotation. Note that

$$\measuredangle[\tilde x^1\,\substack{\tilde x^3\\ \dot x^2}] = \min\{\, \measuredangle[\tilde x^1\,\substack{\tilde x^2\\ \tilde p^3}] + \measuredangle[\tilde x^1\,\substack{\tilde p^2\\ \tilde x^3}],\ 2\cdot\pi - (\measuredangle[\tilde x^1\,\substack{\tilde x^2\\ \tilde p^3}] + \measuredangle[\tilde x^1\,\substack{\tilde p^2\\ \tilde x^3}])\,\}.$$

By (ii), the triangles $[\tilde p^3\tilde x^1\tilde x^2]$ and $[\tilde p^2\tilde x^3\tilde x^1]$ do not overlap if and only if

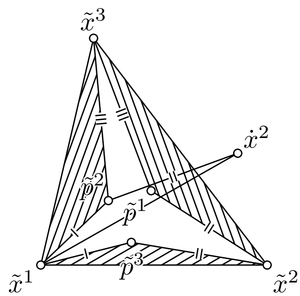

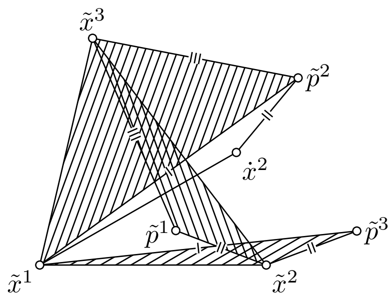

❶ $$2\cdot\pi > \measuredangle[\tilde x^1\,\substack{\tilde x^2\\ \tilde p^3}] + \measuredangle[\tilde x^1\,\substack{\tilde p^2\\ \tilde x^3}] + \measuredangle[\tilde x^1\,\substack{\tilde x^2\\ \tilde x^3}]$$

and

❷ $$\measuredangle[\tilde x^1\,\substack{\tilde x^3\\ \tilde x^2}] > \measuredangle[\tilde x^1\,\substack{\tilde x^3\\ \dot x^2}].$$

The condition ❷ holds if and only if $|\tilde x^2 - \tilde x^3| > |\dot x^2 - \tilde x^3|$, which in turn holds if and only if

❸ $$\begin{aligned}\measuredangle\tilde p^1 &> \measuredangle[\tilde p^2\,\substack{\tilde x^3\\ \dot x^2}]\\ &= \min\{\measuredangle\tilde p^3 + \measuredangle\tilde p^2, 2\cdot\pi - (\measuredangle\tilde p^3 + \measuredangle\tilde p^2)\}.\end{aligned}$$

The inequality follows since the corresponding hinges have the same pairs of sidelengths. (The two pictures show that both possibilities for the minimum can occur.)

Now assume $\measuredangle\tilde p^1 + \measuredangle\tilde p^2 + \measuredangle\tilde p^3 \leqslant 2\cdot\pi$. Then ❸ implies

$$\measuredangle\tilde p^i > \measuredangle\tilde p^j + \measuredangle\tilde p^k.$$

Since no pair of triangles overlap, the same holds for any permutation (i, j, k) of $(1, 2, 3)$. Therefore

$$\measuredangle\tilde p^1 + \measuredangle\tilde p^2 + \measuredangle\tilde p^3 > 2\cdot(\measuredangle\tilde p^1 + \measuredangle\tilde p^2 + \measuredangle\tilde p^3),$$

a contradiction. □

Proof of 10.1; "if" part. Assume $\mathcal{L}$ is geodesic. Consider $x^1, x^2, x^3 \in \in \mathcal{L}$ such that the model triangle $[\tilde{x}^1\tilde{x}^2\tilde{x}^3] = \tilde{\triangle}^\kappa(x^1x^2x^3)$ is defined. Choose $p \in\ ]x^1x^2[$. Applying the short map extension property with $V_3 = \{x^1, x^2, x^3\}$, $V_4 = \{x^1, x^2, x^3, p\}$ and the map $f(x^i) = \tilde{x}^i$, we obtain the point-on-side comparison (8.14*b*).

In case $\mathcal{L}$ is not geodesic, pass to its ultrapower $\mathcal{L}^\varpi$. Note that the short map extension property survives for $\mathcal{L}^\varpi$ and recall that $\mathcal{L}^\varpi$ is geodesic (see 3.5). Thus, from above, $\mathcal{L}^\varpi$ is a complete length CBB(κ) space. By Proposition 8.5, $\mathcal{L}$ is a complete length CBB(κ) space.

"Only if" part. Assume the contrary: $\mathcal{L}$ is complete and CBB(κ), and $x^1, x^2, x^3, p \in \mathcal{L}$ and $\tilde{x}^1, \tilde{x}^2, \tilde{x}^3 \in \mathbb{M}^2(\kappa)$ are such that $|\tilde{x}^i - \tilde{x}^j| \leqslant |x^i - x^j|$ for all i, j but there is no point $\tilde{p} \in \mathbb{M}^2(\kappa)$ such that $|\tilde{p} - \tilde{x}^i| \leqslant |p - x^i|$ for all i.

Note that in this case all comparison triangles $\tilde{\triangle}^\kappa(px^ix^j)$ are defined. This is always true if $\kappa \leqslant 0$. If $\kappa > 0$, and say $\tilde{\triangle}^\kappa(px^1x^2)$ is undefined, then

$$\begin{aligned}|p - x^1| + |p - x^2| &\geqslant 2\cdot\varpi^\kappa - |x^1 - x^2| \geqslant \\ &\geqslant 2\cdot\varpi^\kappa - |\tilde{x}^1 - \tilde{x}^2| \geqslant \\ &\geqslant |\tilde{x}^1 - \tilde{x}^3| + |\tilde{x}^2 - \tilde{x}^3|.\end{aligned}$$

Then the last inequality must be an equality. Thus we may extend by taking $\tilde{p}$ on $[\tilde{x}^1\tilde{x}^3]$ or $[\tilde{x}^2\tilde{x}^3]$.

For each $i \in \{1, 2, 3\}$, consider a point $\tilde{p}^i \in \mathbb{M}^2(\kappa)$ such that $|\tilde{p}^i - \tilde{x}^i|$ is minimal among points satisfying $|\tilde{p}^i - \tilde{x}^j| \leqslant |p - x^j|$ for all $j \neq i$. Clearly, every $\tilde{p}^i$ is inside the triangle $[\tilde{x}^1\tilde{x}^2\tilde{x}^3]$ (that is, in $\mathrm{Conv}(\tilde{x}^1, \tilde{x}^2, \tilde{x}^3)$), and $|\tilde{p}^i - \tilde{x}^i| > |p - x^i|$ for each i. Since the function $x \mapsto \tilde{\measuredangle}^\kappa\{x; a, b\}$ is increasing, it follows that

(i) $|\tilde{p}^i - \tilde{x}^j| = |p - x^j|$ for $i \neq j$;
(ii) no pair of triangles from $[\tilde{p}^1\tilde{x}^2\tilde{x}^3]$, $[\tilde{p}^2\tilde{x}^3\tilde{x}^1]$, $[\tilde{p}^3\tilde{x}^1\tilde{x}^2]$ overlap in $[\tilde{x}^1\tilde{x}^2\tilde{x}^3]$.

As follows from the overlap lemma (10.2), in this case

$$\measuredangle[\tilde{p}^1\,{}^{\tilde{x}^2}_{\tilde{x}^3}] + \measuredangle[\tilde{p}^2\,{}^{\tilde{x}^3}_{\tilde{x}^1}] + \measuredangle[\tilde{p}^3\,{}^{\tilde{x}^1}_{\tilde{x}^2}] > 2\cdot\pi.$$

Since $|\tilde{x}^i - \tilde{x}^j| \leqslant |x^i - x^j|$ we have

$$\measuredangle[\tilde{p}^k\,{}^{\tilde{x}^i}_{\tilde{x}^j}] \leqslant \tilde{\measuredangle}^\kappa(p\,{}^{x^i}_{x^j})$$

if (i, j, k) is a permutation of $(1, 2, 3)$. Therefore

$$\tilde{\measuredangle}^\kappa(p\,{}^{x^1}_{x^2}) + \tilde{\measuredangle}^\kappa(p\,{}^{x^2}_{x^3}) + \tilde{\measuredangle}^\kappa(p\,{}^{x^3}_{x^1}) > 2\cdot\pi,$$

contradicting the CBB(κ) comparison (8.2). □

10.3. Theorem. *Assume any pair of points at distance* $< \varpi^\kappa$ *in the metric space* $\mathcal{U}$ *are joined by a unique geodesic. Then* $\mathcal{U}$ *is* CAT(κ) *if and only if for any 3-point set* V_3 *with perimeter* $< 2\cdot\varpi^\kappa$ *and any 4-point set* $V_4 \supset V_3$ *in* $\mathbb{M}^2(\kappa)$*, any short map* $f\colon V_3 \to \mathcal{U}$ *can be extended to a short map* $F\colon V_4 \to \mathcal{U}$*.*

Note that the "only if" part of Theorem 10.3 does not follow directly from Kirszbraun's theorem, since the desired extension is in $\mathcal{U}$ — not its completion.

10.4. Lemma. *Let* $x^1, x^2, x^3, y^1, y^2, y^3 \in \mathbb{M}(\kappa)$ *be points such that* $|x^i - - x^j| \geqslant |y^i - y^j|$ *for all* i, j*. Then there is a short map* $\Phi\colon \mathbb{M}(\kappa) \to \mathbb{M}(\kappa)$ *such that* $\Phi(x^i) = y^i$ *for all* i*; moreover, one can choose* Φ *so that*

$$\operatorname{Im}\Phi \subset \operatorname{Conv}(y^1, y^2, y^3).$$

We only give an idea of the proof of this lemma; alternatively, it can be obtained as a corollary of Kirszbraun's theorem (10.14)

Idea of the proof. The map Φ can be constructed as a composition of an isometry of $\mathbb{M}(\kappa)$ and the following folding map: given a halfspace H in $\mathbb{M}(\kappa)$, consider the map $\mathbb{M}(\kappa) \to H$ that is the identity on H and reflects all points outside of H into H. This map is a path isometry; in particular, it is short.

The last part of the lemma can be proved by composing this map with folding maps along the sides of triangle $[y^1y^2y^3]$, and passing to a partial limit. □

Proof of 10.3; "if" part. The point-on-side comparison (9.14*b*) follows by taking $V_3 = \{\tilde{x}, \tilde{y}, \tilde{p}\}$ and $V_4 = \{\tilde{x}, \tilde{y}, \tilde{p}, \tilde{z}\}$ where $z \in \,]xy[$. It is only necessary to observe that $F(\tilde{z}) = z$ by uniqueness of $[xy]$.

"Only if" part. Let $V_3 = \{\tilde{x}^1, \tilde{x}^2, \tilde{x}^3\}$ and $V_4 = \{\tilde{x}^1, \tilde{x}^2, \tilde{x}^3, \tilde{p}\}$.

Set $y^i = f(\tilde{x}^i)$ for all i. We need to find a point $q \in \mathcal{U}$ such that $|y^i - q| \leqslant |\tilde{x}^i - \tilde{p}|$ for all i.

Let D be the convex set in $\mathbb{M}^2(\kappa)$ bounded by the model triangle $[\tilde{y}^1\tilde{y}^2\tilde{y}^3] = \tilde{\triangle}^\kappa y^1y^2y^3$; that is, $D = \operatorname{Conv}(\tilde{y}^1, \tilde{y}^2, \tilde{y}^3)$.

Note that $|\tilde{y}^i - \tilde{y}^j| = |y^i - y^j| \leqslant |\tilde{x}^i - \tilde{x}^j|$ for all i, j. Applying Lemma 10.4, we get a short map $\Phi\colon \mathbb{M}(\kappa) \to D$ such that $\Phi\colon \tilde{x}^i \mapsto \tilde{y}^i$.

Further, by the majorization theorem (9.54), there is a short map $F\colon D \to \mathcal{U}$ such that $\tilde{y}^i \mapsto y^i$ for all i.

Thus one can take $q = F \circ \Phi(\tilde{p})$. □

10.5. Exercise. *Assume* $\mathcal{X}$ *is a complete length space that satisfies the following condition: any 4-point subset admits a distance-preserving map to the Euclidean 3-space.*

Prove that $\mathcal{X}$ is isometric to a closed convex subset of a Hilbert space.

10.6. Exercise. *Let $\mathcal{F}_s$ be the metric on the 5-point set $\{p,q,x,y,z\}$ for which $|p-q|=s$ and all the remaining distances are equal 1. For which values s does the space $\mathcal{F}_s$ admit a distance-preserving map into*

a) a complete length $\mathrm{CAT}(0)$ *space?*

b) a complete length $\mathrm{CBB}(0)$ *space?*

The following exercise describes the first known definition of spaces with curvature bounded below; it was given by Abraham Wald [128].

10.7. Exercise. *Let $\mathcal{L}$ be a metric space and $\kappa\leqslant 0$. Prove that $\mathcal{L}$ is* $\mathrm{CBB}(\kappa)$ *if and only if any quadruple of points $p,q,r,s\in\mathcal{L}$ admits a distance-preserving embedding into some $\mathbb{M}^2(\mathrm{K})$ for some $\mathrm{K}\geqslant\kappa$.*

Is the same true for $\kappa>0$; what is the difference?

B (1+n)-point comparison

The following theorem gives a more sensitive analog of the $\mathrm{CBB}(\kappa)$ comparison (8.2).

10.8. (1+n)-point comparison. *Let $\mathcal{L}$ be a complete length* $\mathrm{CBB}(\kappa)$ *space. Then for any array $(p,x^1,\dots,x^n)$ of points in $\mathcal{L}$ there is a model array $(\tilde p,\tilde x^1,\dots,\tilde x^n)$ in $\mathbb{M}^n(\kappa)$ such that*

a) $|\tilde p-\tilde x^i|=|p-x^i|$ for all i.

b) $|\tilde x^i-\tilde x^j|\geqslant|x^i-x^j|$ for all i,j.

Proof. It is sufficient to show that given $\varepsilon>0$ there is an array $(\tilde p,\tilde x^1,\dots$ $\dots,\tilde x^n)$ in $\mathbb{M}^n(\kappa)$ such that

$$|\tilde x^i-\tilde x^j|\geqslant|x^i-x^j|\quad\text{and}\quad\big||\tilde p-\tilde x^i|-|p-x^i|\big|\leqslant\varepsilon.$$

Then one can pass to a limit array for $\varepsilon\to 0+$.

According to 8.11, the set $\mathrm{Str}(x^1,\dots,x^n)$ is dense in $\mathcal{L}$. Thus there is a point $p'\in\mathrm{Str}(\tilde x^1,\dots,\tilde x^n)$ such that $|p'-p|\leqslant\varepsilon$. According to Corollary 13.37, $\mathrm{T}_{p'}$ contains a subcone E isometric to a Euclidean space and containing all vectors $\log[p'x^i]$. Passing to a subspace if necessary, we may assume that $\dim E\leqslant n$.

Mark a point $\tilde p\in\mathbb{M}^n(\kappa)$ and choose a distance-preserving map $\iota\colon E\to$ $\to\mathrm{T}_{\tilde p}\mathbb{M}^n(\kappa)$. Let

$$\tilde x^i=\exp_{\tilde p}\circ\iota(\log[p'x^i]).$$

Thus $|\tilde p-\tilde x^i|=|p'-x^i|$ and therefore

$$\big||\tilde p-\tilde x^i|-|p-x^i|\big|\leqslant|p-p'|\leqslant\varepsilon.$$

From the hinge comparison (8.14c) we have

$$\tilde\measuredangle^\kappa(\tilde p\,{}^{\tilde x^i}_{\tilde x^j}) = \measuredangle[\tilde p\,{}^{\tilde x^i}_{\tilde x^j}] = \measuredangle[p'\,{}^{x^i}_{x^j}] \geqslant \tilde\measuredangle^\kappa(p'\,{}^{x^i}_{x^j}),$$

and thus

$$|\tilde x^i - \tilde x^j| \geqslant |x^i - x^j|.$$

□

10.9. Exercise. *Let $(p, x_1, \dots, x_n)$ be a point array in a* CBB(0) *space. Consider the $n{\times}n$-matrix M with components*

$$m_{i,j} = \tfrac{1}{2}\cdot(|x_i - p|^2 + |x_j - p|^2 - |x_i - x_j|^2).$$

Show that

❶ $$\boldsymbol{s}\cdot M\cdot\boldsymbol{s}^\top \geqslant 0$$

for any vector $\boldsymbol{s} = (s_1, \dots, s_n)$ with nonnegative components.

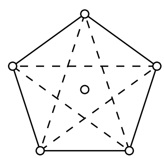

The above exercise describes the so-called Lang–Schroeder–Sturm inequality; it was discovered by Urs Lang and Viktor Schroeder [79] and rediscovered by Karl-Theodor Sturm [124]. It turns out to be weaker than $(1{+}n)$-point comparison. An example can be constructed by perturbing the 6-point metric isometric to a regular pentagon with its center, making its sides slightly longer and diagonals slightly shorter [83]. In particular, this inequality in general metric spaces (not necessarily length spaces) does not imply the inequality in the following exercise.

10.10. Exercise. *Let $\mathcal{L}$ be a complete length* CBB(κ) *space. Show that for any points $p, x_1, x_2, x_3, x_4, x_5$ in $\mathcal{L}$ we have*

$$\tilde\measuredangle^\kappa(p\,{}^{x_1}_{x_5}) + \tilde\measuredangle^\kappa(p\,{}^{x_2}_{x_1}) + \tilde\measuredangle^\kappa(p\,{}^{x_3}_{x_2}) + \tilde\measuredangle^\kappa(p\,{}^{x_4}_{x_3}) + \tilde\measuredangle^\kappa(p\,{}^{x_5}_{x_4}) \leqslant 4\cdot\pi,$$

assuming that the left-hand side is defined.

10.11. Exercise. *Give an example of a metric on a finite set that satisfies the comparison inequality*

$$\tilde\measuredangle^0(p\,{}^{x_1}_{x_2}) + \tilde\measuredangle^0(p\,{}^{x_2}_{x_3}) + \tilde\measuredangle^0(p\,{}^{x_3}_{x_1}) \leqslant 2\cdot\pi$$

for any quadruple of points (p, x_1, x_2, x_3), but is not isometric to a subset of an Alexandrov space with curvature $\geqslant 0$.

10.12. Exercise. *Let $\mathcal{L}$ be a complete length* CBB(κ) *space. Assume that a point array $(a^0, a^1, \dots, a^k)$ in $\mathcal{L}$ is κ-strutting (Definition 15.1) for a point $p \in \mathcal{L}$. Show that there are points $\tilde p, \tilde a^0, \dots, \tilde a^m$ in $\mathbb{M}^{m+1}(\kappa)$ such that*

$$|\tilde p - \tilde a^i| = |p - a^i| \quad \textit{and} \quad |\tilde a^i - \tilde a^j| = |a^i - a^j|$$

for all i and j.

C Helly's theorem

10.13. Helly's theorem. *Let $\mathcal{U}$ be a complete length* CAT(0) *space and $\{K_\alpha\}_{\alpha\in\mathcal{A}}$ be an arbitrary collection of closed bounded convex subsets of $\mathcal{U}$.*

If

$$\bigcap_{\alpha\in\mathcal{A}} K_\alpha = \varnothing,$$

then there is an index array $(\alpha_1, \alpha_2, \dots, \alpha_n)$ in $\mathcal{A}$ such that

$$\bigcap_{i=1}^{n} K_{\alpha_i} = \varnothing.$$

Remarks.

(i) In general, none of the K_α may be compact; otherwise the statement is trivial.

(ii) If $\mathcal{U}$ is a Hilbert space (not necessarily separable), then Helly's theorem is equivalent to the following statement: if a convex bounded set is closed in the ordinary topology then it is compact in the weak topology. One can define weak topology in an arbitrary metric space by taking exteriors of closed ball as prebase. Then Helly's theorem implies the analogous statement for complete length CAT(0) spaces (compare to [96]).

We present the original proof of Urs Lang and Viktor Schroeder [79].

Proof of 10.13. Assume the contrary. Then for any finite set $F \subset \mathcal{A}$,

$$K_F := \bigcap_{\alpha\in F} K_\alpha \neq \varnothing.$$

We will construct a point z such that $z \in K_\alpha$ for each α. Thus we will arrive at a contradiction since

$$\bigcap_{\alpha\in\mathcal{A}} K_\alpha = \varnothing.$$

Choose a point $p \in \mathcal{U}$, and let $r = \sup |K_F - p|$ where F runs over all finite subsets of $\mathcal{A}$. Let p_F^* be the closest point on K_F to p; according to the closest-point projection lemma (9.71), p_F^* exists and is unique.

Take a nested sequence of finite subsets $F_1 \subset F_2 \subset \dots$ of $\mathcal{A}$, such that $|K_{F_n} - p| \to r$.

Let us show that the sequence $p_{F_n}^*$ is Cauchy. Indeed, if not, then for some fixed $\varepsilon > 0$, we can choose two subsequences y_n' and y_n'' of $p_{F_n}^*$

such that $|y'_n - y''_n| \geqslant \varepsilon$. Let z_n be the midpoint of $[y'_n y''_n]$. From the point-on-side comparison (8.14*b*), there is $\delta > 0$ such that

$$|p - z_n| \leqslant \max\{|p - y'_n|, |p - y''_n|\} - \delta.$$

Thus

$$\varlimsup_{n\to\infty} |p - z_n| < r.$$

On the other hand, from convexity, each F_n contains all $z_{\bar{k}}$ with sufficiently large $\bar{k}$, a contradiction.

Thus, $p^*_{F_n}$ converges and we can take $z = \lim_n p^*_{F_n}$. Clearly

$$|p - z| = r.$$

Repeat the above arguments for the sequence $F'_n = F_n \cup \{\alpha\}$. As a result, we get another point z' such that $|p - z| = |p - z'| = r$ and $z, z' \in \in K_{F_n}$ for all n. Thus, if $z \neq z'$ the midpoint $\hat{z}$ of $[zz']$ would belong to all K_{F_n}, and from comparison, we would have $|p - \hat{z}| < r$, a contradiction.

Thus, $z' = z$; in particular $z \in K_\alpha$ for each $\alpha \in \mathcal{A}$. □

D Kirszbraun's theorem

A slightly weaker version of the following theorem was proved by Urs Lang and Viktor Schroeder [79].

10.14. Kirszbraun's theorem. *Let $\mathcal{L}$ be a complete length* CBB(κ) *space, $\mathcal{U}$ be a complete length* CAT(κ) *space, $Q \subset \mathcal{L}$ be arbitrary subset and $f\colon Q \to \mathcal{U}$ be a short map. Assume that there is $z \in \mathcal{U}$ such that $f(Q) \subset \mathrm{B}[z, \frac{\varpi^\kappa}{2}]_{\mathcal{U}}$. Then $f\colon Q \to \mathcal{U}$ can be extended to a short map $F\colon \mathcal{L} \to \mathcal{U}$ (that is, there is a short map $F\colon \mathcal{L} \to \mathcal{U}$ such that $F|_Q = f$).*

The condition $f(Q) \subset \mathrm{B}[z, \frac{\varpi^\kappa}{2}]$ trivially holds for any $\kappa \leqslant 0$ since in this case $\varpi^\kappa = \infty$. The following example shows that this condition is needed for $\kappa > 0$.

Conjecture 10.22 (if true) gives an equivalent condition for the existence of a short extension; it states that the following example is the only obstacle.

10.15. Example. *Let $\mathbb{S}^m_+$ be a closed m-dimensional unit hemisphere. Denote its boundary, which is isometric to $\mathbb{S}^{m-1}$, by $\partial\mathbb{S}^m_+$. Clearly, $\mathbb{S}^m_+$ is* CBB(1) *and $\partial\mathbb{S}^m_+$ is* CAT(1), *but the identity map $\partial\mathbb{S}^m_+ \to \partial\mathbb{S}^m_+$ cannot be extended to a short map $\mathbb{S}^m_+ \to \partial\mathbb{S}^m_+$ (there is no place for the pole).*

There is also a direct generalization of this example to a hemisphere in a Hilbert space of arbitrary cardinal dimension.

First we prove this theorem in the case $\kappa \leqslant 0$ (10.17). In the proof of the more complicated case $\kappa > 0$, we use the case $\kappa = 0$. The following lemma is the main ingredient in the proof.

10.16. Finite+one lemma. *Let $\kappa \leqslant 0$, $\mathcal{L}$ be a complete length* CBB(κ) *space, and $\mathcal{U}$ be a complete length* CAT(κ) *space. Suppose $x^1, x^2, \dots, x^n$ in $\mathcal{L}$ and $y^1, y^2, \dots, y^n$ in $\mathcal{U}$ are such that $|x^i - x^j| \geqslant |y^i - y^j|$ for all i, j.*

Then for any $p \in \mathcal{L}$, there is $q \in \mathcal{U}$ such that $|y^i - q| \leqslant |x^i - p|$ for each i.

Proof. It is sufficient to prove the lemma only for $\kappa = 0$ and -1. The proofs of these two cases are identical, only the formulas differ. In the proof, we assume $\kappa = 0$ and provide the formulas for $\kappa = -1$ in the footnotes.

From the $(1+n)$-point comparison (10.8), there is a model configuration $\tilde{p}, \tilde{x}^1, \tilde{x}^2, \dots, \tilde{x}^n$ in $\mathbb{M}^n(\kappa)$ such that $|\tilde{p} - \tilde{x}^i| = |p - x^i|$ and $|\tilde{x}^i - \tilde{x}^j| \geqslant |x^i - x^j|$ for all i, j. It follows that we can assume that $\mathcal{L} = \mathbb{M}^n(\kappa)$.

For each i, consider functions $f^i\colon \mathcal{U} \to \mathbb{R}$ and $\tilde{f}^i\colon \mathbb{M}^n(\kappa) \to \mathbb{R}$ defined as follows:[1]

$$(A)^0 \qquad f^i = \tfrac{1}{2}\cdot \operatorname{dist}^2_{y^i}, \quad \tilde{f}^i = \tfrac{1}{2}\cdot \operatorname{dist}^2_{\tilde{x}^i}.$$

Consider $\boldsymbol{f} = (f^1, f^2, \dots, f^n)\colon \mathcal{U} \to \mathbb{R}^n$ and $\tilde{\boldsymbol{f}} = (\tilde{f}^1, \tilde{f}^2, \dots \dots, \tilde{f}^n)\colon \mathbb{M}^n(\kappa) \to \mathbb{R}^n$.

Define

$$\begin{aligned} \operatorname{Up} \boldsymbol{f}(\mathcal{U}) &= \{\, \boldsymbol{v} \in \mathbb{R}^{k+1} : \exists\, \boldsymbol{w} \in \boldsymbol{f}(\mathcal{U}) \text{ such that } \boldsymbol{v} \succcurlyeq \boldsymbol{w} \,\}, \\ \operatorname{Min} \boldsymbol{f}(\mathcal{U}) &= \{\, \boldsymbol{v} \in \boldsymbol{f}(\mathcal{U}) : \text{if } \boldsymbol{v} \succcurlyeq \boldsymbol{w} \in \boldsymbol{f}(\mathcal{U}) \text{ then } \boldsymbol{w} = \boldsymbol{v} \,\}. \end{aligned}$$

(See Definition 14.1.) Note it is sufficient to prove that $\tilde{\boldsymbol{f}}(\tilde{p}) \in \operatorname{Up} \boldsymbol{f}(\mathcal{U})$.

Clearly, $(f^i)'' \geqslant 1$. Thus by Theorem 14.3*a*, the set $\operatorname{Up} \boldsymbol{f}(\mathcal{U}) \subset \mathbb{R}^n$ is convex.

Arguing by contradiction, let us assume that $\tilde{\boldsymbol{f}}(\tilde{p}) \notin \operatorname{Up} \boldsymbol{f}(\mathcal{U})$.

Then there exists a supporting hyperplane $\alpha_1 \cdot x_1 + \dots \alpha_n \cdot x_n = c$ to $\operatorname{Up} \boldsymbol{f}(\mathcal{U})$, separating it from $\tilde{\boldsymbol{f}}(\tilde{p})$. According to Lemma 14.6*b*, $\alpha_i \geqslant 0$ for each i. So we may assume that $(\alpha_1, \alpha_2, \dots, \alpha_n) \in \Delta^{n-1}$ (that is, $\alpha_i \geqslant 0$ for each i and $\sum \alpha_i = 1$ and

$$\sum_i \alpha_i \cdot \tilde{f}^i(\tilde{p})) < \inf \left\{ \sum_i \alpha_i \cdot f^i(q) : q \in \mathcal{U} \right\}.$$

---

[1] In case $\kappa = -1$,

$$(A)^- \qquad f^i = \cosh \circ \operatorname{dist}_{y^i}, \quad \tilde{f}^i = \cosh \circ \operatorname{dist}_{\tilde{x}^i}.$$

The latter contradicts the following claim.

❶ *Given* $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \dots, \alpha_n) \in \Delta^{n-1}$, *let*

$$h = \sum_i \alpha_i \cdot f^i \qquad h\colon \mathcal{U} \to \mathbb{R} \qquad z = \operatorname{MinPoint} h \in \mathcal{U}$$
$$\tilde h = \sum_i \alpha_i \cdot \tilde f^i \qquad \tilde h\colon \mathbb{M}^n(\kappa) \to \mathbb{R} \qquad \tilde z = \operatorname{MinPoint} \tilde h \in \mathbb{M}^n(\kappa).$$

Then $h(z) \leqslant \tilde h(\tilde z)$.

Proof of the claim. Note that $\boldsymbol{d}_z h \geqslant 0$. Thus, for each i, we have[2]

$$(B)^0 \qquad \begin{aligned} 0 &\leqslant (\boldsymbol{d}_z h)(\uparrow_{[zy^i]}) = \\ &= -\sum_j \alpha_j \cdot |z - y^j| \cdot \cos \measuredangle[z\,{}^{y^i}_{y^j}] \leqslant \\ &\leqslant -\sum_j \alpha_j \cdot |z - y^j| \cdot \cos \tilde\measuredangle^0(z\,{}^{y^i}_{y^j}) = \\ &= -\tfrac{1}{2\cdot|z-y^i|} \cdot \sum_j \alpha_j \cdot \left[|z-y^i|^2 + |z-y^j|^2 - |y^i - y^j|^2\right]. \end{aligned}$$

In particular[3],

$$(C)^0 \qquad \sum_i \alpha_i \cdot \left[\sum_j \alpha_j \cdot \left[|z-y^i|^2 + |z-y^j|^2 - |y^i - y^j|^2\right]\right] \leqslant 0,$$

or[4]

$$(D)^0 \qquad 2\cdot h(z) \leqslant \sum_{i,j} \alpha_i \cdot \alpha_j \cdot |y^i - y^j|^2.$$

---

[2] In case $\kappa = -1$, the same calculations give

$$(B)^- \qquad 0 \leqslant \dots \leqslant -\tfrac{1}{\sinh|z-y^i|} \cdot \sum_j \alpha_j \cdot \left[\cosh|z-y^i| \cdot \cosh|z-y^j| - \cosh|y^i - y^j|\right].$$

[3] In case $\kappa = -1$, the same calculations give

$$(C)^- \qquad \sum_i \alpha_i \cdot \left[\sum_j \alpha_j \cdot \left[\cosh|z-y^i| \cdot \cosh|z-y^j| - \cosh|y^i - y^j|\right]\right] \leqslant 0.$$

[4] In case $\kappa = -1$,

$$(D)^- \qquad (h(z))^2 \leqslant \sum_{i,j} \alpha_i \cdot \alpha_j \cdot \cosh|y^i - y^j|.$$

Note that if $\mathcal{U} \overset{iso}{=\!=} \mathbb{M}^n(\kappa)$, then all inequalities in (B, C, D) are sharp. Thus the same argument as above, repeated for $\tilde{x}^1, \tilde{x}^2, \dots, \tilde{x}^n$ in $\mathbb{M}^n(\kappa)$, gives[5]

$$(E)^0 \qquad 2\cdot\tilde{h}(\tilde{z}) = \sum_{i,j} \alpha_i\cdot\alpha_j\cdot|\tilde{x}^i - \tilde{x}^j|^2.$$

Note that

$$|\tilde{x}^i - \tilde{x}^j| \geqslant |x^i - x^j| \geqslant |y^i - y^j|$$

for all i, j. Thus, (D) and (E) imply the claim. △□

10.17. Kirszbraun's theorem for nonpositive bound. *Let $\kappa \leqslant 0$, $\mathcal{L}$ be a complete length* CBB(κ) *space, $\mathcal{U}$ be a complete length* CAT(κ) *space, $Q \subset \mathcal{L}$ be arbitrary subset and $f\colon Q \to \mathcal{U}$ be a short map. Then there is a short extension $F\colon \mathcal{L} \to \mathcal{U}$ of f; that is, there is a short map $F\colon \mathcal{L} \to \mathcal{U}$ such that $F|_Q = f$.*

Remark. If $\mathcal{U}$ is proper, then we do not need Helly's theorem (10.13); compactness of closed balls in $\mathcal{U}$ is sufficient in this case.

Proof of 10.17. By Zorn's lemma, we can assume that $Q \subset \mathcal{L}$ is a maximal set; that is, $f\colon Q \to \mathcal{U}$ does not admit a short extension to any larger set $Q' \supset Q$.

Let us argue by contradiction. Assume that $Q \neq \mathcal{L}$ and choose $p \in \in \mathcal{L}\backslash Q$. Then

$$\bigcap_{x\in Q} \overline{\mathrm{B}}[f(x), |p - x|] = \varnothing.$$

Since $\kappa \leqslant 0$, the balls are convex; thus, by Helly's theorem (10.13), one can choose points $x^1, x^2, \dots, x^n$ in Q such that

$$❷ \qquad \bigcap_{i=1}^{n} \overline{\mathrm{B}}[y^i, |x^i - p|] = \varnothing,$$

where $y^i = f(x^i)$. Finally note that ❷ contradicts the finite+one lemma (10.16). □

Proof of Kirszbraun's theorem (10.14). The case $\kappa \leqslant 0$ is already proved in 10.17. Thus it remains to prove the theorem only in case $\kappa > 0$. After rescaling we can assume that $\kappa = 1$ and therefore $\varpi^\kappa = \pi$.

[5] In case $\kappa = -1$,

$$(E)^- \qquad (\tilde{h}(\tilde{z}))^2 = \sum_{i,j} \alpha_i\cdot\alpha_j\cdot\cosh|\tilde{x}^i - \tilde{x}^j|.$$

Since $\overline{\mathrm{B}}[z,\pi/2]_{\mathcal{U}}$ is a complete length CAT(κ) space, we can assume $\mathcal{U} = \overline{\mathrm{B}}[z,\pi/2]_{\mathcal{U}}$. In particular, $\operatorname{diam}\mathcal{U} \leqslant \pi$.

Further, any two points $x, y \in \mathcal{U}$ such that $|x - y| < \pi$ are joined by a unique geodesic; if $|x - y| = \pi$, then the concatenation of $[xz]$ and $[zy]$ as a geodesic from x to y. Hence $\mathcal{U}$ is geodesic.

We may also assume that $\operatorname{diam}\mathcal{L} \leqslant \pi$. Otherwise $\mathcal{L}$ is one-dimensional (see 8.43); in this case the result follows since $\mathcal{U}$ is geodesic.

Assume the theorem is false. Then there is a set $Q \subset \mathcal{L}$, a short map $f\colon Q \to \mathcal{U}$, and $p \in \mathcal{L}\backslash Q$ such that

❸ $$\bigcap_{x\in Q} \overline{\mathrm{B}}[f(x), |x - p|] = \varnothing.$$

We will apply 10.17 for $\kappa = 0$ to the Euclidean cones $\mathring{\mathcal{L}} = \operatorname{Cone}\mathcal{L}$ and $\mathring{\mathcal{U}} = \operatorname{Cone}\mathcal{U}$. Note that

- ⋄ $\mathring{\mathcal{U}}$ is a complete length CAT(0) space (see 11.7a),
- ⋄ since $\operatorname{diam}\mathcal{L} \leqslant \pi$ we have $\mathring{\mathcal{L}}$ is CBB(0) (see 11.6a).

Further, we will view the spaces $\mathcal{L}$ and $\mathcal{U}$ as unit spheres in $\mathring{\mathcal{L}}$ and $\mathring{\mathcal{U}}$ respectively. In the cones $\mathring{\mathcal{L}}$ and $\mathring{\mathcal{U}}$ we will use "$|{*}|$" for distance to the vertex, denoted by 0, "$\cdot$" for cone multiplication, "$\measuredangle(x,y)$" for $\measuredangle[0\,{}^{x}_{y}]$, and "$\langle x,y\rangle$" for $|x|\cdot|y|\cdot\cos\measuredangle[0\,{}^{x}_{y}]$. In particular,

- ⋄ $|x - y|_{\mathcal{L}} = \measuredangle(x,y)$ for any $x, y \in \mathcal{L}$,
- ⋄ $|x - y|_{\mathcal{U}} = \measuredangle(x,y)$ for any $x, y \in \mathcal{U}$,
- ⋄ for any $y \in \mathcal{U}$, we have

❹ $$\measuredangle(z,y) \leqslant \tfrac{\pi}{2}.$$

Let $\mathring{Q} = \operatorname{Cone} Q \subset \mathring{\mathcal{L}}$ and let $\mathring{f}\colon \mathring{Q} \to \mathring{\mathcal{U}}$ be the natural cone extension of f; that is, $y = f(x) \Rightarrow t\cdot y = \mathring{f}(t\cdot x)$ for $t \geqslant 0$. Clearly $\mathring{f}$ is short.

Applying 10.17 for $\mathring{f}$, we get a short extension map $\mathring{F}\colon \mathring{\mathcal{L}} \to \mathring{\mathcal{U}}$. Let $s = \mathring{F}(p)$. Then

❺ $$|s - \mathring{f}(w)| \leqslant |p - w|$$

for any $w \in \mathring{Q}$. In particular, $|s| \leqslant 1$. Applying ❺ for $w = t\cdot x$ and $t \to \infty$ we have

❻ $$\langle f(x), s\rangle \geqslant \cos\measuredangle(p,x)$$

for any $x \in Q$.

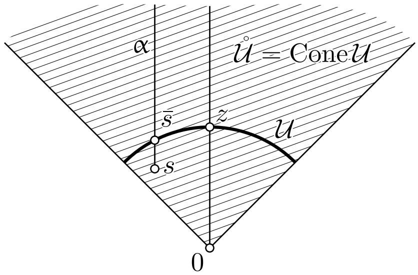

By comparison, the geodesics $\mathrm{geod}_{[s\ t\cdot z]}$ converge as $t \to \to \infty$ and the limit is a half-line; denote it by $\alpha\colon [0,\infty) \to \to \mathring{\mathcal{U}}$. From ❹, the function $t \mapsto \langle f(x), \alpha(t)\rangle$ is nondecreasing. From ❻, for the necessarily unique point $\bar{s}$ on the half-line α such that $|\bar{s}| = 1$, we also have

$$\langle f(x), \bar{s}\rangle \geqslant \cos\measuredangle(p, x)$$

or

$$\measuredangle(\bar{s}, f(x)) \leqslant \measuredangle(p, f(x))$$

for any $x \in Q$, in contradiction to ❸. □

10.18. Exercise. *Let $\mathcal{U}$ be* CAT(0)*. Assume there are two point arrays, $(x^0, x^1, \dots, x^{\bar{k}})$ in $\mathcal{U}$ and $(\tilde{x}^0, \tilde{x}^1, \dots, \tilde{x}^{\bar{k}})$ in $\mathbb{E}^m$, such that $|x^i - x^j|_{\mathcal{U}} = |\tilde{x}^i - \tilde{x}^j|_{\mathbb{E}^m}$ for each i, j, and for any point $z_0 \in \mathcal{U}$ there is $i > 0$ such that $|z_0 - x_i| \geqslant |x_0 - x_i|$.*

Prove that there is a subset $Q \subset \mathcal{L}$ isometric to a convex set in $\mathbb{E}^m$ and containing all the points x^i.

10.19. Exercise. *Let $\mathcal{L}$ be a complete length* CBB(0) *space, and $(x^0, x^1, \dots, x^{\bar{k}})$ in $\mathcal{L}$ and $(\tilde{x}^0, \tilde{x}^1, \dots, \tilde{x}^{\bar{k}})$ in $\mathbb{E}^m$ be two point arrays such that $|x^i - x^j|_{\mathcal{L}} = |\tilde{x}^i - \tilde{x}^j|_{\mathbb{E}^m}$ for each i, j. Assume $\tilde{x}^0$ lies in the interior of* $\mathrm{Conv}(\tilde{x}^1, \dots, \tilde{x}^{\bar{k}})$.

Prove that there is a subset $Q \subset \mathcal{L}$ isometric to a convex set in $\mathbb{E}^m$ and containing all the points x^i.

The following statement we call $(2n+2)$-point comparison.

10.20. Exercise. *Let $\mathcal{U}$ be a complete length* CAT(κ) *space. Consider $x, y \in \mathcal{U}$ and an array $((p^1, q^1), (p^2, q^2), \dots, (p^n, q^n))$ of pairs of points in $\mathcal{U}$, such that there is a model configuration $\tilde{x}$, $\tilde{y}$ and array of pairs $((\tilde{p}^1, \tilde{q}^1), (\tilde{p}^2, \tilde{q}^2), \dots, (\tilde{p}^n, \tilde{q}^n))$ in $\mathbb{M}^3(\kappa)$ with the following properties:*

a) $[\tilde{x}\tilde{p}^1\tilde{q}^1] = \tilde{\triangle}^{\kappa} x p^1 q^1$ and $[\tilde{y}\tilde{p}^n\tilde{q}^n] = \tilde{\triangle}^{\kappa} y p^n q^n$,

b) the simplex $\tilde{p}^i\tilde{p}^{i+1}\tilde{q}^i\tilde{q}^{i+1}$ is a model simplex[6] of $p^i p^{i+1} q^i q^{i+1}$ for all i.

Then for any choice of n points $\tilde{z}^i \in [\tilde{p}^i\tilde{q}^i]$, we have

$$|\tilde{x} - \tilde{z}^1| + |\tilde{z}^1 - \tilde{z}^2| + \ldots + |\tilde{z}^{n-1} - \tilde{z}^n| + |\tilde{z}^n - \tilde{y}| \geqslant |x - y|.$$

[6]that is, $|\tilde{p}^i - \tilde{q}^i| = |p^i - q^i|$, $|\tilde{p}^i - \tilde{p}^{i+1}| = |p^i - p^{i+1}|$, $|\tilde{q}^i - \tilde{q}^{i+1}| = |q^i - q^{i+1}|$, $|\tilde{p}^i - \tilde{q}^{i+1}| = |p^i - q^{i+1}|$ and $|\tilde{p}^{i+1} - \tilde{q}^i| = |p^{i+1} - q^i|$.

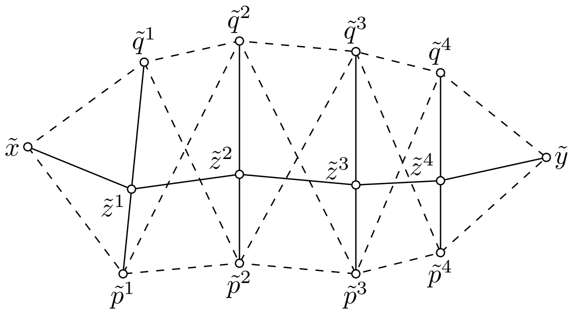

E Remarks and open problems

10.21. Open problem. *Find a necessary and sufficient condition for a finite metric space to admit distance-preserving embeddings into*

a) some length CBB(κ) *space,*

b) some length CAT(κ) *space.*

A metric on a finite set $\{a^1, a^2, \dots, a^n\}$, can be described by the matrix with components

$$s^{ij} = |a^i - a^j|^2,$$

which we will call the associated matrix. The set of associated matrices of all metrics that admit a distance-preserving map into a CBB(0) or a CAT(0) space form a convex cone. The latter follows since the rescalings and products of CBB(0) (or CAT(0)) spaces are CBB(0) (or CAT(0) respectively). This convexity gives hope that the cone admits an explicit description.

For the 5-point CAT(0) case, the (2+2)-comparison is a necessary and sufficient condition. This was proved by Tetsu Toyoda [126]; another proof was found by Nina Lebedeva and the third author [85]. For the 5-point CBB(0) case, the (1+4)-comparison is a necessary and sufficient condition; it was proved by Nina Lebedeva and the third author [84]. Starting from the 6-point case, only some necessary and some sufficient conditions are known; for more on the subject see [6, 83, 85].

The following conjecture (if true) would give the right generality for Kirszbraun's theorem (10.14). It states that the example 10.15 is the only obstacle to extending short maps.

10.22. Conjecture. *Assume $\mathcal{L}$ is a complete length* CBB(1) *space, $\mathcal{U}$ is a complete length* CAT(1) *space, $Q \subset \mathcal{L}$ is a proper subset, and $f\colon Q \to \mathcal{U}$ is a short map that does not admit a short extension to any bigger set $Q' \supset Q$. Then:*

a) Q is isometric to a sphere in a Hilbert space (of finite or cardinal dimension). Moreover, there is a point $p \in \mathcal{L}$ such that $|p - q| = \frac{\pi}{2}$ for any $q \in Q$.

b) The map $f\colon Q \to \mathcal{U}$ is a distance-preserving map and there is no point $p' \in \mathcal{U}$ such that $|p' - q'| = \frac{\pi}{2}$ for any $q' \in f(Q)$.

Curvature-free analogs. Let us present a collection of exercises on curvature-free analogs of Kirszbraun's theorem. It is worthwhile to know these results despite they are far from Alexandrov geometry.

The following exercise gives an analog of the finite+one lemma (10.16), discovered by the third author and Stephan Stadler [104].

10.23. Exercise. *Let $\mathcal{X}$ and $\mathcal{Y}$ be metric spaces, $A \subset \mathcal{X}$, and $f\colon A \to \mathcal{Y}$ be a short map. Assume $\mathcal{Y}$ is compact and for any finite set $F \subset \mathcal{X}$ there is a short map $F \to \mathcal{Y}$ that agrees with f on $F \cap A$. Then there is a short map $\mathcal{X} \to \mathcal{Y}$ that agrees with f on A.*

The following statement was first observed by John Isbell [70].

10.24. Exercise. *We say that a metric space $\mathcal{X}$ is injective if for an arbitrary metric space $\mathcal{Z}$ and a subset $Q \subset \mathcal{Z}$, any short map $Q \to \mathcal{X}$ can be extended as a short map $\mathcal{Z} \to \mathcal{X}$.*

a) Prove that any metric space $\mathcal{X}$ admits a distance-preserving embedding into an injective metric space.

b) Use this to construct an analog of convex hull in the category of metric spaces; this is called the injective hull.

10.25. Exercise. *We say that a compact space $\mathcal{X}$ is a Kirszbraun source if for an arbitrary complete metric space $\mathcal{Z}$ and subset $Q \subset \mathcal{X}$, any short map $Q \to \mathcal{Z}$ can be extended to a short map $\mathcal{X} \to \mathcal{Z}$.*

Prove that a metric space $\mathcal{X}$ is a Kirszbraun source if and only if it satisfies the ultratriangle inequality for all triples of points; that is,

$$|x - z| \leqslant \max\{|x - y|, |y - z|\}$$

for any $x, y, z \in \mathcal{X}$.

Chapter 11

Warped products

The warped product is a construction that produces a new metric space, denoted by $\mathcal{B} \times_f \mathcal{F}$, from two metric spaces $\mathcal{B}$, $\mathcal{F}$ and a function $f\colon \mathcal{B} \to \mathbb{R}_{\geqslant 0}$.

Many important constructions such as direct product, cone, spherical suspension, and join can be defined using warped products.

A Definition

First we define the warped product for length spaces and then we expand the definition to allow for arbitrary metric spaces $\mathcal{F}$.

Let $\mathcal{B}$ and $\mathcal{F}$ be length spaces and $f\colon \mathcal{B} \to [0,\infty)$ be a continuous function.

For any path $\gamma\colon [0,1] \to \mathcal{B} \times \mathcal{F}$, we write $\gamma = (\gamma_{\mathcal{B}}, \gamma_{\mathcal{F}})$ where $\gamma_{\mathcal{B}}$ is the projection of γ to $\mathcal{B}$, and $\gamma_{\mathcal{F}}$ is the projection to $\mathcal{F}$. If $\gamma_{\mathcal{B}}$ and $\gamma_{\mathcal{F}}$ are Lipschitz, set

❶ $$\operatorname{length}_f \gamma := \int_0^1 \sqrt{v_{\mathcal{B}}^2 + (f\circ\gamma_{\mathcal{B}})^2\cdot v_{\mathcal{F}}^2}\cdot dt,$$

where $\int$ is Lebesgue integral, and $v_{\mathcal{B}}$ and $v_{\mathcal{F}}$ are the speeds of $\gamma_{\mathcal{B}}$ and $\gamma_{\mathcal{F}}$ respectively.

Note that $\operatorname{length}_f \gamma \geqslant \operatorname{length} \gamma_{\mathcal{B}}$. The integral in ❶ can be broken into the sum of two parts: one for the restriction of $\operatorname{length}_f \gamma$ as in ❶ to the nonzero set of $f\circ\gamma_{\mathcal{B}}$, and the other for the length of $\gamma_{\mathcal{B}}$ restricted to the zero set.

Consider the pseudometric on $\mathcal{B} \times_f \mathcal{F}$ defined by

$$|x-y| := \inf\left\{\operatorname{length}_f \gamma : \gamma(0) = x, \gamma(1) = y\right\}$$

where the exact lower bound is taken for all Lipschitz paths $\gamma\colon[0,1]\to\to\mathcal{B}\times\mathcal{F}$. The corresponding metric space is called the warped product with base $\mathcal{B}$, fiber $\mathcal{F}$ and warping function f; it will be denoted by $\mathcal{B}\times_f\mathcal{F}$.

The points in $\mathcal{B}\times_f\mathcal{F}$ can be described by corresponding pairs $(p,\varphi)\in\in\mathcal{B}\times\mathcal{F}$. Note that if $f(p)=0$ for some $p\in\mathcal{B}$, then $(p,\varphi)=(p,\psi)$ for any $\varphi,\psi\in\mathcal{F}$.

We do not claim that every Lipschitz curve in $\mathcal{B}\times_f\mathcal{F}$ may be reparametrized as the image of a Lipschitz curve in $\mathcal{B}\times\mathcal{F}$; in fact this is not true.

11.1. Proposition. *The warped product $\mathcal{B}\times_f\mathcal{F}$ satisfies:*

a) *The projection $(p,\varphi_0)\mapsto p$ is a submetry which when restricted to any horizontal leaf $\mathcal{B}\times\{\varphi_0\}$ is an isometry to $\mathcal{B}$.*
b) *If $f(p_0)\neq 0$, the projection $(p_0,\varphi)\mapsto\varphi$ of the vertical leaf $\{p_0\}\times\mathcal{F}$, with its length metric, is a homothety onto $\mathcal{F}$ with multiplier $\frac{1}{f(p_0)}$.*
c) *If f achieves its (local) minimum at p_0, then the inclusion of the vertical leaf $\{p_0\}\times\mathcal{F}$ in $\mathcal{B}\times_f\mathcal{F}$ is (locally) distance-preserving.*

Proof. Claim (b) follows from the f-length formula ➊.

Also by ➊, the projection of $\mathcal{B}\times_f\mathcal{F}$ onto $\mathcal{B}\times\{\varphi_0\}$ given by $(p,\varphi)\mapsto\mapsto(p,\varphi_0)$ is length-nonincreasing; hence (a).

The projection $(p,\varphi)\mapsto(p_0,\varphi)$ of a neighborhood of the vertical leaf $\{p_0\}\times\mathcal{F}$ to $\{p_0\}\times\mathcal{F}$ is length-nonincreasing if p_0 is a local minimum point of f. If p_0 is a global minimum point of f, then the same holds for the projection of whole space. Hence (c). □

Note that any horizontal leaf is weakly convex, but does not have to be convex even if $\mathcal{B}\times_f\mathcal{F}$ is a geodesic space, since vanishing of the warping function f allows geodesics to bifurcate into distinct horizontal leaves. For instance, if there is a geodesic with the ends in the zero set

$$Z=\{\,(p,\varphi)\in\mathcal{B}\times_f\mathcal{F})\,:\,f(p)=0\,\},$$

then there is a geodesic with the same ends in each horizontal leaf.

11.2. Proposition. *Suppose $\mathcal{B}$ and $\mathcal{F}$ are length spaces and $f\colon\mathcal{B}\to\to[0,\infty)$ is a continuous function. Then the warped product $\mathcal{B}\times_f\mathcal{F}$ is a length space.*

Proof. It is sufficient to show that for any $\alpha\colon[0,1]\to\mathcal{B}\times_f\mathcal{F}$ there is a path $\beta\colon[0,1]\to\mathcal{B}\times\mathcal{F}$ with the same endpoints such that

$$\operatorname{length}\alpha\geqslant\operatorname{length}_f\beta.$$

If $f\circ\alpha_{\mathcal{B}}(t)>0$ for any t, then the vertical projection $\alpha_{\mathcal{F}}$ is defined. In this case, let $\beta(t)=(\alpha_{\mathcal{B}}(t),\alpha_{\mathcal{F}}(t))\in\mathcal{B}\times\mathcal{F}$. Clearly

$$\operatorname{length}\alpha=\operatorname{length}_f\beta.$$

If $f\circ\alpha_{\mathcal{B}}(t_0)=0$ for some t_0, let β be the concatenation of three curves in $\mathcal{B}\times\mathcal{F}$, namely: (1) the horizontal curve $(\alpha_{\mathcal{B}}(t),\varphi)$ for $t\leqslant t_0$, (2) a vertical path in form (s,φ) to (s,ψ) and (3) the horizontal curve $(\alpha_{\mathcal{B}}(t),\psi)$ for $t\geqslant t_0$. By ❶, the f-length of the middle path in the concatenation is vanishing; therefore the f-length of α cannot be smaller than length of $\alpha_{\mathcal{B}}$, that is,

$$\operatorname{length}_f\alpha\geqslant\operatorname{length}\alpha_{\mathcal{B}}=\operatorname{length}_f\beta.$$

The statement follows. □

Distance in a warped product is fiber-independent, in the sense that distances may be calculated by substituting for $\mathcal{F}$ a different length space:

11.3. Fiber-independence theorem. *Consider length spaces $\mathcal{B}$, $\mathcal{F}$ and $\check{\mathcal{F}}$, and a locally Lipschitz function $f\colon\mathcal{B}\to\mathbb{R}_{\geqslant0}$. Assume $p,q\in\mathcal{B}$, $\varphi,\psi\in\mathcal{F}$ and $\check\varphi,\check\psi\in\check{\mathcal{F}}$. Then*

$$|\varphi-\psi|_{\mathcal{F}}\geqslant|\check\varphi-\check\psi|_{\check{\mathcal{F}}}$$
$$\Downarrow$$
$$|(p,\varphi)-(q,\psi)|_{\mathcal{B}\times_f\mathcal{F}}\geqslant|(p,\check\varphi)-(q,\check\psi)|_{\mathcal{B}\times_f\check{\mathcal{F}}}.$$

In particular,

$$|(p,\varphi)-(q,\psi)|_{\mathcal{B}\times_f\mathcal{F}}=|(p,0)-(q,\ell)|_{\mathcal{B}\times_f\mathbb{R}},$$

where $\ell=|\varphi-\psi|_{\mathcal{F}}$.

Proof. Let γ be a path in $(\mathcal{B}\times\mathcal{F})$.

Since $|\varphi-\psi|_{\mathcal{F}}\geqslant|\check\varphi-\check\psi|_{\check{\mathcal{F}}}$, there is a Lipschitz path $\gamma_{\check{\mathcal{F}}}$ from $\check\varphi$ to $\check\psi$ in $\check{\mathcal{F}}$ such that

$$(\operatorname{speed}\gamma_{\mathcal{F}})(t)\geqslant(\operatorname{speed}\gamma_{\check{\mathcal{F}}})(t)$$

for almost all $t\in[0,1]$. Consider the path $\check\gamma=(\gamma_{\mathcal{B}},\gamma_{\check{\mathcal{F}}})$ from $(p,\check\varphi)$ to $(q,\check\psi)$ in $\mathcal{B}\times_f\check{\mathcal{F}}$. Clearly

$$\operatorname{length}_f\gamma\geqslant\operatorname{length}_f\check\gamma.$$

□

11.4. Exercise. *Let $\mathcal{B}$ and $\mathcal{F}$ be length spaces and $f,g\colon\mathcal{B}\to\mathbb{R}_{\geqslant}$ be two locally Lipschitz nonnegative functions. Assume $f(b)\leqslant g(b)$ for any $b\in\mathcal{B}$. Show that $\mathcal{B}\times_f\mathcal{F}\leqslant\mathcal{B}\times_g\mathcal{F}$; that is, there is a distance-noncontracting map $\mathcal{B}\times_f\mathcal{F}\to\mathcal{B}\times_g\mathcal{F}$.*

B Extended definitions

The fiber-independence theorem implies that

$$|(p,\varphi)-(q,\psi)|_{\mathcal{B}\times_f\mathcal{F}}=|(p,0)-(q,|\varphi-\psi|_{\mathcal{F}})|_{\mathcal{B}\times_f\mathbb{R}}$$

for any $(p,\varphi),(q,\psi)\in\mathcal{B}\times\mathcal{F}$. In particular, if $\iota\colon A\to\check{A}$ is an isometry between two subsets $A\subset\mathcal{F}$ and $\check{A}\subset\check{\mathcal{F}}$ in length spaces $\mathcal{F}$ and $\check{\mathcal{F}}$, and $\mathcal{B}$ is a length space, then for any warping function $f\colon\mathcal{B}\to\mathbb{R}_{\geqslant 0}$, the map ι induces an isometry between the sets $\mathcal{B}\times_f A\subset\mathcal{B}\times_f\mathcal{F}$ and $\mathcal{B}\times_f\check{A}\subset$ $\subset\mathcal{B}\times_f\check{\mathcal{F}}$.

This observation allows us to define the warped product $\mathcal{B}\times_f\mathcal{F}$ where the fiber $\mathcal{F}$ does not carry its length metric. Indeed we can use Kuratowsky embedding to realize $\mathcal{F}$ as a subspace in a length space, say $\mathcal{F}'$. Therefore we can take the warped product $\mathcal{B}\times_f\mathcal{F}'$ and identify $\mathcal{B}\times_f\mathcal{F}$ with its subspace consisting of all pairs (b,φ) such that $\varphi\in\mathcal{F}$. According to the Fiber-independence theorem 11.3, the resulting space does not depend on the choice of $\mathcal{F}'$.

C Examples

Direct product. The simplest example is the direct product $\mathcal{B}\times\mathcal{F}$, which can also be written as the warped product $\mathcal{B}\times_1\mathcal{F}$. That is, for $p,q\in\mathcal{B}$ and $\varphi,\psi\in\mathcal{F}$, the latter metric simplifies to

$$|(p,\varphi)-(q,\psi)|=\sqrt{|p-q|^2+|\varphi-\psi|^2}.$$

Cones. The Euclidean cone $\operatorname{Cone}\mathcal{F}$ over a metric space $\mathcal{F}$ can be written as the warped product $[0,\infty)\times_{\mathrm{id}}\mathcal{F}$. That is, for $s,t\in[0,\infty)$ and $\varphi,\psi\in\mathcal{F}$, the metric is given by the cosine rule

$$|(s,\varphi)-(t,\psi)|=\sqrt{s^2+t^2-2{\cdot}s{\cdot}t{\cdot}\cos\alpha},$$

where $\alpha=\min\{\pi,|\varphi-\psi|\}$. (See Section 6E.)

Instead of the Euclidean cosine rule, we may use the cosine rule in $\mathbb{M}^2(\kappa)$:

$$|(s,\varphi)-(t,\psi)|=\tilde{\gamma}^{\kappa}\{\alpha;s,t\}.$$

In this way we get κ-cones over $\mathcal{F}$, denoted by $\operatorname{Cone}^\kappa\mathcal{F}=[0,\infty)\times_{\operatorname{sn}^\kappa}\mathcal{F}$ for $\kappa\leqslant 0$ and $\operatorname{Cone}^\kappa\mathcal{F}=[0,\varpi^\kappa]\times_{\operatorname{sn}^\kappa}\mathcal{F}$ for $\kappa>0$.

The 1-cone $\operatorname{Cone}^1\mathcal{F}$ is also called the spherical suspension over $\mathcal{F}$; it is also denoted by $\operatorname{Susp}\mathcal{F}$. That is,

$$\operatorname{Susp}\mathcal{F}=[0,\pi]\times_{\sin}\mathcal{F}.$$

11.5. Exercise. *Let $\mathcal{F}$ be a length space and $A \subset \mathcal{F}$. Show that* $\mathrm{Cone}^\kappa A$ *is convex in* $\mathrm{Cone}^\kappa \mathcal{F}$ *if and only if A is π-convex in $\mathcal{F}$.*

Doubling. The doubling space $\mathcal{W}$ of a metric space $\mathcal{V}$ on a closed subset $A \subset \mathcal{V}$ can be also defined as a special type of warped product. Consider the fiber $\mathbb{S}^0$ consisting of two points with distance 2 from each other. Then $\mathcal{W}$ is isometric to the warped product with base $\mathcal{V}$, fiber $\mathbb{S}^0$ and warping function dist_A; that is,

$$\mathcal{W} \stackrel{iso}{=\!=} \mathcal{V} \times_{\mathrm{dist}_A} \mathbb{S}^0.$$

D 1-dimensional base

The following theorems provide conditions for the spaces and functions in a warped product with 1-dimensional base to have curvature bounds. These theorems are originally due to Valerii Berestovskii [17]. They are baby cases of the characterization of curvature bounds in warped products given in [4, 5].

11.6. Theorem.

a) If $\mathcal{L}$ is a complete length $\mathrm{CBB}(1)$ *space and* $\operatorname{diam}\mathcal{L} \leqslant \pi$, *then*

$$\begin{aligned} \mathrm{Susp}\,\mathcal{L} &= [0,\pi] \times_{\sin} \mathcal{L} \quad \textit{is} \quad \mathrm{CBB}(1), \\ \mathrm{Cone}\,\mathcal{L} &= [0,\infty) \times_{\mathrm{id}} \mathcal{L} \quad \textit{is} \quad \mathrm{CBB}(0), \\ \mathrm{Cone}^{-1}\,\mathcal{L} &= [0,\infty) \times_{\sinh} \mathcal{L} \quad \textit{is} \quad \mathrm{CBB}(-1). \end{aligned}$$

Moreover, the converse also holds in each of the three cases.

b) If $\mathcal{L}$ is a complete length $\mathrm{CBB}(0)$ *space, then*

$$\begin{aligned} \mathbb{R} \times \mathcal{L} \quad &\textit{is a complete length } \mathrm{CBB}(0) \textit{ space}, \\ \mathbb{R} \times_{\exp} \mathcal{L} \quad &\textit{is a complete length } \mathrm{CBB}(-1) \textit{ space}. \end{aligned}$$

Moreover, the converse also holds in each of the two cases.

c) If $\mathcal{L}$ is a complete length $\mathrm{CBB}(-1)$ *space, then* $\mathbb{R} \times_{\cosh} \mathcal{L}$ *is a complete length* $\mathrm{CBB}(-1)$ *space. Moreover, the converse also holds.*

11.7. Theorem. *Let $\mathcal{L}$ be a metric space.*

a) If $\mathcal{L}$ is $\mathrm{CAT}(1)$, *then*

$$\begin{aligned} \mathrm{Susp}\,\mathcal{L} &= [0,\pi] \times_{\sin} \mathcal{L} \quad \textit{is} \quad \mathrm{CAT}(1), \\ \mathrm{Cone}\,\mathcal{L} &= [0,\infty) \times_{\mathrm{id}} \mathcal{L} \quad \textit{is} \quad \mathrm{CAT}(0), \\ \mathrm{Cone}^{-1}\,\mathcal{L} &= [0,\infty) \times_{\sinh} \mathcal{L} \quad \textit{is} \quad \mathrm{CAT}(-1). \end{aligned}$$

Moreover, the converse also holds in each of the three cases.

b) If $\mathcal{L}$ *is a complete length* CAT(0) *space, then* $\mathbb{R}\times\mathcal{L}$ *is* CAT(0) *and* $\mathbb{R}\times_{\exp}\mathcal{L}$ *is* CAT(−1)*. Moreover, the converse also holds in each of the two cases.*

c) If $\mathcal{L}$ *is* CAT(−1) *then* $\mathbb{R}\times_{\cosh}\mathcal{L}$ *is* CAT(−1)*. Moreover, the converse also holds.*

In the proof of the above two theorems we will use the following proposition.

11.8. Proposition.

a)

$$\begin{aligned}\operatorname{Susp}\mathbb{S}^{m-1} &= [0,\pi]\times_{\sin}\mathbb{S}^{m-1} \stackrel{iso}{=\!=} \mathbb{S}^m,\\ \operatorname{Cone}\mathbb{S}^{m-1} &= [0,\infty)\times_{\mathrm{id}}\mathbb{S}^{m-1} \stackrel{iso}{=\!=} \mathbb{E}^m,\\ \operatorname{Cone}^{-1}\mathbb{S}^{m-1} &= [0,\infty)\times_{\sinh}\mathbb{S}^{m-1} \stackrel{iso}{=\!=} \mathbb{M}^m(-1).\end{aligned}$$

b)

$$\begin{aligned}\mathbb{R}\times\mathbb{E}^{m-1} &\stackrel{iso}{=\!=} \mathbb{E}^m,\\ \mathbb{R}\times_{\exp}\mathbb{E}^{m-1} &\stackrel{iso}{=\!=} \mathbb{M}^m(-1).\end{aligned}$$

c)

$$\mathbb{R}\times_{\cosh}\mathbb{M}^{m-1}(-1) \stackrel{iso}{=\!=} \mathbb{M}^m(-1).$$

The proof is left to the reader.

Proof of 11.6. Let us prove the last statement in (a); the remaining statements are similar. Each proof is based on the fiber-independence theorem 11.3 and the corresponding statement in Proposition 11.8.

Choose an arbitrary quadruple of points

$$(s,\varphi),(t^1,\varphi^1),(t^2,\varphi^2),(t^3,\varphi^3)\in[0,\infty)\times_{\sinh}\mathcal{L}.$$

Since $\operatorname{diam}\mathcal{L}\leqslant\pi$, the $(1{+}n)$-point comparison (10.8) provides a quadruple of points $\psi,\psi^1,\psi^2,\psi^3\in\mathbb{S}^3$ such that

$$|\psi-\psi^i|_{\mathbb{S}^3}=|\varphi-\varphi^i|_{\mathcal{L}}$$

and

$$|\psi^i-\psi^j|_{\mathbb{S}^3}\geqslant|\varphi^i-\varphi^j|_{\mathcal{L}}$$

for all i and j.

According to Proposition 11.8a,

$$\operatorname{Cone}^{-1}\mathbb{S}^3=[0,\infty)\times_{\sinh}\mathbb{S}^3 \stackrel{iso}{=\!=} \mathbb{M}^4(-1).$$

Consider the quadruple of points

$$(s,\psi),(t^1,\psi^1),(t^2,\psi^2),(t^3,\psi^3)\in \operatorname{Cone}^{-1}\mathbb{S}^3=\mathbb{M}^4(-1).$$

By the fiber-independence theorem 11.3,

$$|(s,\psi)-(t^i,\psi^i)|_{[0,\infty)\times_{\sinh}\mathbb{S}^3}=|(s,\varphi)-(t^i,\varphi^i)|_{[0,\infty)\times_{\sinh}\mathcal{L}}$$

and

$$|(t^i,\psi^i)-(t^j,\psi^j)|_{[0,\infty)\times_{\sinh}\mathbb{S}^3}\geqslant|(t^i,\varphi^i)-(t^j,\varphi^j)|_{[0,\infty)\times_{\sinh}\mathcal{L}}$$

for all i and j. Since four points of $\mathbb{M}^4(-1)$ lie in an isometric copy of $\mathbb{M}^3(-1)$, it remains to apply Exercise 8.3. □

11.9. Exercise. *The spherical join $\mathcal{U}\star\mathcal{V}$ of two metric spaces $\mathcal{U}$ and $\mathcal{V}$ is defined as the unit sphere equipped with the angle metric in the product of Euclidean cones* $\operatorname{Cone}\mathcal{U}\times\operatorname{Cone}\mathcal{V}$.

Assume $\mathcal{U}$ and $\mathcal{V}$ are nonempty spaces.

a) Show that $\mathcal{U}\star\mathcal{V}$ is CAT(1) *if and only if $\mathcal{U}$ and $\mathcal{V}$ are* CAT(1).

b) Show that $\mathcal{U}\star\mathcal{V}$ is CBB(1) *if and only if $\mathcal{U}$ and $\mathcal{V}$ are* CBB(1).

E Remarks

Let us formulate general results on curvature bounds of warped products proved by the first author and Richard Bishop [5].

11.10. Theorem. *Let $\mathcal{B}$ be a complete finite-dimensional* CBB(κ) *length space, and the locally Lipschitz function $f\colon\mathcal{B}\to\mathbb{R}_{\geqslant}$ satisfy $f''+\kappa\cdot f\leqslant 0$. If $\mathcal{B}^{(\dagger)}$ is the result of gluing two copies of $\mathcal{B}$ on the closure of the set of boundary points where f is nonvanishing, and $f^\dagger\colon\mathcal{B}^{(\dagger)}\to\mathbb{R}_{\geqslant}$ is the tautological extension of f, suppose $\mathcal{B}^{(\dagger)}$ is* CBB(κ) *and $f^\dagger$ satisfies $(f^\dagger)''+\kappa\cdot f^\dagger\leqslant 0$. Denote by $Z\subset\mathcal{B}$ the zero set of the restriction of f to the boundary of $\mathcal{B}$. Suppose $\mathcal{F}$ is a complete finite-dimensional* CBB(κ′) *space. Then the warped product $\mathcal{B}\times_f\mathcal{F}$ is* CBB(κ) *in the following two cases:*

a) If $Z=\varnothing$ and

$$\kappa'\geqslant\kappa\cdot f^2(b)$$

for any $b\in\mathcal{B}$.

b) If $Z\neq\varnothing$ and

$$|d_zf|^2\leqslant\kappa'$$

for any $z\in Z$.

We mention that in the setting of this theorem, f necessarily vanishes only at boundary points if f is not identically 0.

11.11. Theorem. *Let $\mathcal{B}$ be a complete* $\mathrm{CAT}(\kappa)$ *length space, and the function $f\colon\mathcal{B}\to\mathbb{R}_{\geqslant}$ satisfy $f''+\kappa\cdot f\geqslant 0$, where f is Lipschitz on bounded sets or B is locally compact. Denote by $Z\subset\mathcal{B}$ the zero set of f. Suppose $\mathcal{F}$ is a complete* $\mathrm{CAT}(\kappa')$ *space. Then the warped product $\mathcal{B}\times_f\mathcal{F}$ is* $\mathrm{CAT}(\kappa)$ *in the following two cases:*

a) If $Z=\varnothing$ and

$$\kappa'\leqslant\kappa\cdot f^2(b)$$

for any $b\in\mathcal{B}$.

b) If $Z\neq\varnothing$,

$$[(d_zf)\!\uparrow_{[zb]}]^2\geqslant\kappa'$$

for any minimizing geodesic $[zb]$ from Z to a point $b\in\mathcal{B}$ and

$$\kappa'\leqslant\kappa\cdot f^2(b)$$

for any $b\in\mathcal{B}$ such that $\operatorname{dist}_X b\geqslant\frac{\varpi^\kappa}{2}$.

These theorems will be proved in the next volume.

Chapter 12

Polyhedral spaces

A Definitions

12.1. Definition. *A length space $\mathcal{P}$ is called a piecewise $\mathbb{M}(\kappa)$ polyhedral space if it admits a finite triangulation τ such that an arbitrary simplex σ in τ is isometric to a simplex in the model space $\mathbb{M}^{\dim\sigma}(\kappa)$.*

By triangulation of a piecewise $\mathbb{M}(\kappa)$ polyhedral space we will understand a triangulation as in the definition. If we do not wish to specify κ, we will say that $\mathcal{P}$ is a polyhedral space.

By rescaling we can assume that $\kappa = 1$, 0, or -1.

a) The $\mathbb{M}(1)$-polyhedral spaces will also be called spherical polyhedral spaces;
b) The $\mathbb{M}(0)$-polyhedral spaces will also be called Euclidean polyhedral spaces;
c) The $\mathbb{M}(-1)$-polyhedral spaces will also be called hyperbolic polyhedral spaces.

Note that according to the above definition, all polyhedral spaces are compact. However, most of the statements below admit straightforward generalizations to locally polyhedral space; that is, complete length spaces, any point of which admits a closed neighborhood isometric to a polyhedral space. The latter class of spaces includes in particular infinite covers of polyhedral spaces.

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

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

The simplexes that contain σ form an abstract simplicial complex called the link of σ, denoted by Link_σ. If $m = \dim\sigma$, then the set of vertexes of Link_σ is formed by the $(m+1)$-simplexes that contain σ; the set of its edges are formed by the $(m+2)$-simplexes that contain σ, and so on.

The link Link_σ can be identified with the subcomplex of τ formed by all the simplexes 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 the interior of σ. 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 τ be a triangulation of a polyhedral space $\mathcal{P}$. If a point $p\in\mathcal{P}$ lies in the interior of a $\hbar$-simplex σ of τ then the tangent space $\mathrm{T}_p\mathcal{P}$ is naturally isometric to

$$\mathbb{E}^{\hbar}\times(\mathrm{Cone}\,\mathrm{Link}_\sigma).$$

Equivalently, the space of directions Σ_p can be isometrically identified with the $\hbar$-th spherical suspension over Link_σ; that is,

$$\Sigma_p\stackrel{iso}{=\!=}\mathrm{Susp}^{\hbar}(\mathrm{Link}_\sigma).$$

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 the interior of a simplex σ, the isometry class of Link_σ and $\hbar=\dim\sigma$ determine the isometry class of Σ_p and the other way around.

A small neighborhood of p is isometric to a neighborhood of the tip of the κ-cone over Σ_p. In fact, if this property holds at any point of a compact length space $\mathcal{P}$ then $\mathcal{P}$ is a piecewise $\mathbb{M}(\kappa)$ space [82].

B Curvature bounds

Recall that the definition of ℓ-simply connected space is given in 9.62.

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

12.2. Theorem. *Let $\mathcal{P}$ be a piecewise $\mathbb{M}(\kappa)$ space and τ be a triangulation of $\mathcal{P}$. Then*

a) $\mathcal{P}$ is locally $\mathrm{CAT}(\kappa)$ if and only if any connected component of the link of any simplex σ in τ is $(2\cdot\pi)$-simply connected. Equivalently, if and only if any closed local geodesic in Link_σ has length at least $2\cdot\pi$.

b) *$\mathcal{P}$ is a complete length* CAT(κ) *space if and only if $\mathcal{P}$ is $(2\cdot\varpi^\kappa)$-simply connected and any connected component of the link of any simplex σ in τ is $(2\cdot\pi)$-simply connected.*

Proof. We will prove only the "if" part; the "only if" part is evident by the generalized Hadamard–Cartan theorem (9.63) and Theorem 11.6.

Let us apply induction on $\dim\mathcal{P}$. The base case $\dim\mathcal{P} = 0$ is evident.

Induction 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 tip in the κ-cone over Σ_p. By Theorem 11.6*a*, it is sufficient to show that

❶ Σ_p *is* CAT(1).

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 locally CAT(1). Applying the generalized Hadamard–Cartan theorem (9.63), we get ❶.

To prove (*b*) apply the generalized Hadamard–Cartan theorem to $\mathcal{P}$. □

A metric graph is a finite graph equipped with a length-metric, such that every edge is isometric to a line segment; in other words, a metric graph is a 1-dimension polyhedral space.

12.3. Exercise. *Show that any metric tree is* CAT(κ) *for any* κ.

12.4. Exercise. *Show that if in a Euclidean polyhedral space $\mathcal{P}$ any two points can be connected by a unique geodesic, then $\mathcal{P}$ is* CAT(0).

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

12.5. Theorem. *Let $\mathcal{P}$ be a piecewise $\mathbb{M}(\kappa)$ space and τ be a triangulation of $\mathcal{P}$. Then $\mathcal{P}$ is* CBB(κ) *if and only if the following conditions hold.*

a) *τ is pure; that is, any simplex in τ is a face of some simplex of dimension exactly m.*

b) *The link of any simplex of dimension $m-1$ is formed by a single point or two points.*

c) *Any link of any simplex of dimension $m-2$ has diameter at most π.*

d) *The link of any simplex of dimension $\leqslant m-2$ is connected.*

Remarks. Condition (*d*) can be reformulated in the following way:

d′) Any path $\gamma\colon[0,1]\to\mathcal{P}$ can be approximated by paths $\gamma_n\colon[0,1]\to\mathcal{P}$ that cross only simplexes of dimension m and $m-1$.

Further, modulo the other conditions, condition (c) is equivalent to the following:

c′) The link of any simplex of dimension $m-2$ is isometric to a circle of length $\leqslant 2{\cdot}\pi$ or a closed real interval of length $\leqslant\pi$.

Proof. We will prove the "if" part. The "only if" part is similar and is left to the reader.

We apply induction on m. The base case $m=1$ follows from the assumption (b).

Step. Assume that the theorem is proved for polyhedral spaces of dimesnion less than m. Suppose $\dim\mathcal{P}=m$.

According to the globalization theorem (8.30), it is sufficient to show that $\mathcal{P}$ is locally CBB(κ).

Fix $p\in\mathcal{P}$. Note that a spherical neighborhood of p is isometric to a spherical neighborhood of the tip of the tangent κ-cone

$$\mathrm{T}_p^\kappa=\mathrm{Cone}^\kappa(\Sigma_p).$$

Hence it is sufficient to show that

❷ T_p^κ *is* CBB(κ) *for any* $p\in\mathcal{P}$.

By Theorem 11.6*a*, the latter is equivalent to

❸ $\operatorname{diam}\Sigma_p\leqslant\pi$ *and* Σ_p *is* CBB(1).

If $m=2$, then ❸ follows from (b).

To prove the case $m\geqslant 3$, note that Σ_p is an $(m-1)$-dimensional spherical polyhedral space and all the conditions of the theorem hold for Σ_p. It remains to apply the induction hypothesis. □

12.6. Exercise. *Assume $\mathcal{P}$ is a piecewise $\mathbb{M}(\kappa)$ space and $\dim\mathcal{P}\geqslant 2$. Show that*

a) if $\mathcal{P}$ is CBB(κ')*, then $\kappa'\leqslant\kappa$ and $\mathcal{P}$ is* CBB(κ)*,*

b) if $\mathcal{P}$ is CAT(κ')*, then $\kappa'\geqslant\kappa$ and $\mathcal{P}$ is* CAT(κ)*.*

C Flag complexes

12.7. Definition. *A simplicial complex $\mathcal{S}$ is called flag if whenever $\{v_0,\dots,v_{\Bbbk}\}$ is a set of distinct vertexes of $\mathcal{S}$ that are pairwise joined by edges, then the vertexes $v_0,\dots,v_{\Bbbk}$ span a $\Bbbk$-simplex in $\mathcal{S}$.*

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

Note that every flag complex is determined by its 1-skeleton.

12.8. Exercise. *Show that the barycentric subdivision of any simplicial complex is a flag complex. Conclude that any finite simplicial complex is homeomorphic to a compact length* $\mathrm{CAT}(1)$ *space.*

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

From the definition of flag complex we get the following:

12.10. Lemma. *Any link of a flag complex is flag.*

Proof of Proposition 12.9. By Lemma 12.10, the no-triangle condition holds for any flag complex and all its links.

Now assume 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 vertexes. Repeating this observation for triangles, 4-simplexes, 5-simplexes and so on we get that $\mathcal{S}$ is flag. ☐

Right-angled triangulation. A triangulation of a spherical polyhedral space is called a right-angled 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 a right-angled 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 $\mathrm{CAT}(1)$ spaces glued from right-angled simplexes admit the following characterization discovered by Michael Gromov [61, p. 122].

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

Proof; "only-if" part. Assume there are three vertexes v_1, v_2 and v_3 of τ that are pairwise joined by edges but do not span a simplex. Note that in this case

$$\measuredangle[v_1\,{}^{v_2}_{v_3}] = \measuredangle[v_2\,{}^{v_3}_{v_1}] = \measuredangle[v_3\,{}^{v_1}_{v_1}] = \pi.$$

Equivalently,

❶ *The concatenation of the geodesics $[v_1v_2]$, $[v_2v_3]$ and $[v_3v_1]$ forms a locally geodesic loop in $\mathcal{P}$.*

Now assume that $\mathcal{P}$ is CAT(1). Then by 11.6*a*, $\mathrm{Link}_\sigma \mathcal{P}$ is a compact length CAT(1) space for every simplex σ in τ.

Each of these links is a right-angled spherical complex and by Theorem 12.2, none of these links can contain a geodesic circle of length less than $2\cdot\pi$.

Therefore Proposition 12.9 and ❶ imply the "only-if" part.

"If" part. By Lemma 12.10 and Theorem 12.2, it is sufficient to show that any closed local geodesic γ in a flag complex $\mathcal{S}$ with right-angled metric has length at least $2\cdot\pi$.

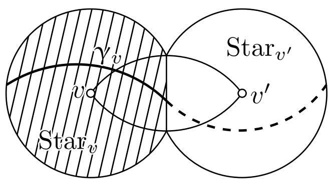

Fix a flag complex $\mathcal{S}$. Recall that the star of a vertex v (briefly $\overline{\mathrm{Star}}_v$) is formed by all the simplexes containing v. Similarly, Star_v, the open star of v, is the union of all simplexes containing v with faces opposite v removed.

Choose a simplex σ that contains a point of γ. Let v be a vertex of σ. Set $f(t) = \cos|v-\gamma(t)|$. Note that

$$f''(t) + f(t) = 0$$

if $f(t) > 0$. Since the zeroes of f are π apart, γ spends time π on every visit to Star_v.

After leaving Star_v, the local geodesic γ must enter another simplex, say σ', which has a vertex v' not joined to v by an edge.

Since τ is flag, the open stars Star_v and $\mathrm{Star}_{v'}$ do not overlap. 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 $2\cdot\pi$. □

12.12. Exercise. *Let p be a point in a product of metric trees. Show that a closed geodesic in the space of directions Σ_p has length either $2\cdot\pi$ or at least $3\cdot\pi$.*

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

12.14. Exercise. *Let $\mathcal{U}$ be a complete length* CAT(0) *space and $\varphi_1, \varphi_2, \dots, \varphi_k : \mathcal{U} \to \mathcal{U}$ be commuting short retractions; that is,*

- ◇ $\varphi_i \circ \varphi_i = \varphi_i$ *for each i;*
- ◇ $\varphi_i \circ \varphi_j = \varphi_j \circ \varphi_i$ *for any i and j;*
- ◇ $|\varphi_i(x) - \varphi_i(y)|_{\mathcal{U}} \leqslant |x-y|_{\mathcal{U}}$ *for each i and any $x, y \in \mathcal{U}$.*

Set $A_i = \mathrm{Im}\,\varphi_i$ for all i. Note that each A_i is a weakly convex set.

Assume Γ is a finite graph (without loops and multiple edges) with edges labeled by $1, 2, \dots, n$. Denote by $\mathcal{U}^\Gamma$ the space obtained by taking a *copy of $\mathcal{U}$ for each vertex of Γ and gluing two such copies along A_i if the corresponding vertexes are joined by an edge labeled by i.*

Show that $\mathcal{U}^\Gamma$ is CAT(0)

The space of trees. The following construction is given by Louis Billera, Susan Holmes, and Karen Vogtmann [21].

Let $\mathcal{T}_n$ be the set of all metric trees with n end-vertexes labeled by $a_1, \dots, a_n$. To describe one tree in $\mathcal{T}_n$ we may fix a topological tree τ with end vertexes $a_1, \dots, a_n$ and all the other vertexes of degree 3, and prescribe the lengths of $2{\cdot}n - 3$ edges. If the length of an edge is 0, we assume that edge degenerates; such a tree can be also described using a different topological tree τ'. The subset of $\mathcal{T}_n$ corresponding to the given topological tree τ can be identified with a convex closed cone in $\mathbb{R}^{2\cdot n-3}$. 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.

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

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

Formally, a cubical complex is defined as a subcomplex of the unit cube in Euclidean space of large dimension; that is, a collection of faces of the cube, that with each face 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, that is, contains a finite number of cubes.

The union of all the cubes in a cubical complex $\mathcal{Q}$ will be called its underlying space; it will be denoted by $\mathcal{Q}$ or by $\underline{\mathcal{Q}}$ if we need to emphasize that we are talking about a set, not a complex. A homeomorphism from $\underline{\mathcal{Q}}$ 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 same dimension.

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

The link of each cube in a cubical complex admits a natural right-angled triangulation; each simplex corresponds to an adjusted cube.

12.16. Exercise. *Show that a cubical complex $\mathcal{Q}$ is locally* CAT(0) *if and only if the link of each vertex in $\mathcal{Q}$ is flag.*

D Remarks

The condition on polyhedral CAT(κ) spaces given in Theorem 12.2 might look easy to use, but in fact, it is hard to check even in very simple cases. For example the description of those coverings of $\mathbb{S}^3$ that branch at three great circles and are CAT(1) requires quite a bit of work; an answer is given by Ruth Charney and Michael Davis [38].

Analogs of the flag condition for spherical Coxeter simplexes could resolve the following problem.

12.17. Braid space problem. *Consider $\mathbb{C}^n$ with coordinates $z_1, \dots$ $\dots, z_n$. Let us remove from $\mathbb{C}^n$ the complex hyperplanes $z_i = z_j$ for all $i \neq j$, pass to the universal cover, and consider the completion $\mathcal{B}_n$ of the obtained space.*

Is it true that $\mathcal{B}_n$ is CAT(0) *for any n?*

The above question has an affirmative answer for $n \leqslant 3$ and is open for all $n \geqslant 4$ [38, 100].

Recall that by the Hadamard–Cartan theorem (9.61), any complete length CAT(0) space is contractible. Therefore any complete length locally CAT(0) space is aspherical; that is, has contractible universal cover.

This observation can be used together with Exercise 12.16 to construct examples of exotic aspherical spaces; for example, compact topological manifolds with universal cover not homeomorphic to a Euclidean space. A survey on the subject is given by Michael Davis [45]; a more elementary introduction to the subject is given by the authors [8, Chapter 3].

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 [39].

The CAT(0) property of a cube complex admits interesting (and useful) geometric descriptions if one replaces the ℓ^2-metric with a natural ℓ^1- or ℓ^∞-metric on each cube. The following statement was proved by Brian Bowditch [24].

12.18. 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; that is, for any metric space $\mathcal{X}$ with a subset A, any short map $A \to (Q, \ell^\infty)$ can be extended to a short map $\mathcal{X} \to (Q, \ell^\infty)$.

c) A cube complex Q equipped with ℓ^1-metric is median; that is, for any three points x, y, z there is a unique point m (called the median of x, y, and z) that lies on some geodesics $[xy]$, $[xz]$ and $[yz]$.

Part III

Structure and tools

Chapter 13

First order differentiation

A Ultratangent space

The following theorem is often used together with the observation that the ultralimit of any sequence of length spaces is geodesic (see 3.6).

13.1. Theorem.

a) If $\mathcal{L}$ is a complete length CBB(κ) *space and $p \in \mathcal{L}$, then T_p^ω is* CBB(0).

b) If $\mathcal{U}$ is a complete length CAT(κ) *space and $p \in \mathcal{U}$, then T_p^ω is* CAT(0).

The proofs of both parts are nearly identical.

Proof; (a). Since $\mathcal{L}$ is a complete length CBB(κ) space, then its blowup $n{\cdot}\mathcal{L}$ (see Section 6H) is a complete length CBB(κ/n^2) space. By Proposition 8.5, the ω-blowup $\omega{\cdot}\mathcal{L}$ is CBB(0) and so is T_p^ω as a metric component of $\omega{\cdot}\mathcal{L}$.

(b). Since $\mathcal{U}$ is a complete length CAT(κ) space, then its blowup $n{\cdot}\mathcal{U}$ is CAT(κ/n^2). By Proposition 9.7, $\omega{\cdot}\mathcal{U}$ is CAT(0) and so is T_p^ω as a metric component of $\omega{\cdot}\mathcal{U}$. □

Recall that the tangent space T_p can be considered as a subset of T_p^ω (see 6.15). Therefore we have the following:

13.2. Corollary.

a) If $\mathcal{L}$ is a complete length CBB(κ) *space and $p \in \mathcal{L}$, then T_p is* CBB(0). *Moreover, T_p satisfies the (1+n)-point comparison (10.8).*

b) If $\mathcal{U}$ is a complete length CAT(κ) *space and $p \in \mathcal{U}$, then T_p is* CAT(0). *Moreover, T_p satisfies the (2n+2)-comparison (10.20).*

13.3. Proposition. *Assume $\mathcal{Z}$ is a complete length* CBB *or* CAT *space and $f\colon \mathcal{Z} \hookrightarrow \mathbb{R}$ is a semiconcave locally Lipscitz subfunction. Then for any $p \in \operatorname{Dom} f$, the ultradifferential $\boldsymbol{d}_p^{\omega}\colon \mathrm{T}_p^{\omega} \to \mathbb{R}$ is a concave function.*

Proof. Fix a geodesic $[x^{\omega}y^{\omega}]$ in T_p^{ω}.

It is sufficient to show that for any subarc $[\bar{x}^{\omega}\bar{y}^{\omega}]$ of $[x^{\omega}y^{\omega}]$ that does not contain the ends there is a sequence of geodesics $[\bar{x}^n\bar{y}^n]$ in $n{\cdot}\mathcal{Z}$ converging to $[\bar{x}^{\omega}\bar{y}^{\omega}]$.

Choose any sequences $\bar{x}^n, \bar{y}^n \in n{\cdot}\mathcal{Z}$ such that $\bar{x}^n \to \bar{x}^{\omega}$ and $\bar{y}^n \to \bar{y}^{\omega}$ as $n \to \omega$. Note that $[\bar{x}^n\bar{y}^n]$ converges to $[\bar{x}^{\omega}\bar{y}^{\omega}]$ as $n \to \omega$. The latter holds trivially in the CAT case, and the CBB case follows from 8.37. □

B Length property of tangent space

13.4. Theorem. *Let $\mathcal{U}$ be a complete length* CAT(κ) *space and $p \in \mathcal{U}$. Then $\mathrm{T}_p\mathcal{U}$ is a length space.*

This theorem together with 13.2 imply the following.

13.5. Corollary. *Let $\mathcal{U}$ be a complete length* CAT(κ) *space and $p \in \mathcal{U}$. Then $\mathrm{T}_p\mathcal{U}$ is a complete length* CAT(0) *space.*

Proof of Theorem 13.4. Since $\mathrm{T}_p = \operatorname{Cone}\Sigma_p$, it is sufficient to show that for any hinge $[p\,{}^x_y]$ such that $\measuredangle[p\,{}^x_y] < \pi$ and any $\varepsilon > 0$, there is $z \in \mathcal{U}$ such that

❶ $$\measuredangle[p\,{}^x_z] < \tfrac{1}{2}{\cdot}\measuredangle[p\,{}^x_y] + \varepsilon, \quad \measuredangle[p\,{}^y_z] < \tfrac{1}{2}{\cdot}\measuredangle[p\,{}^x_y] + \varepsilon.$$

Fix a small positive number $\delta \ll \varepsilon$. Let $\bar{x} \in \,]px]$ and $\bar{y} \in \,]py]$ denote the points such that $|p-\bar{x}| = |p-\bar{y}| = \delta$. Let z denote the midpoint between $\bar{x}$ and $\bar{y}$.

Since $\delta \ll \varepsilon$ we can assume that

$$\tilde{\measuredangle}^{\kappa}(p\,{}^{\bar{x}}_{\bar{y}}) < \measuredangle[p\,{}^x_y] + \varepsilon.$$

By Alexandrov's lemma (6.2), we have

$$\tilde{\measuredangle}^{\kappa}(p\,{}^{\bar{x}}_{z}) + \tilde{\measuredangle}^{\kappa}(p\,{}^{\bar{y}}_{z}) < \tilde{\measuredangle}^{\kappa}(p\,{}^{\bar{x}}_{\bar{y}}).$$

By construction,

$$\tilde{\measuredangle}^{\kappa}(p\,{}^{\bar{x}}_{z}) = \tilde{\measuredangle}^{\kappa}(p\,{}^{\bar{y}}_{z}).$$

Applying the angle comparison (9.14c), we get ❶. □

The following example was constructed by Stephanie Halbeisen [65]. It shows that an analogous statement does not hold for CBB spaces. If the dimension is finite, such examples do not exist; for proper spaces the question is open, see 13.38.

13.6. Example. *There is a complete length* CBB *space* $\check{\mathcal{L}}$ *with a point* $p \in \check{\mathcal{L}}$ *such that the space of directions* $\Sigma_p\check{\mathcal{L}}$ *is not a* π*-length space, and therefore the tangent space* $\mathrm{T}_p\check{\mathcal{L}}$ *is not a length space.*

Construction. Let $\mathbb{H}$ be a Hilbert space formed by infinite sequences of real numbers $\boldsymbol{x} = (x_0, x_1, \ldots)$ with norm $|\boldsymbol{x}|^2 = \sum_i (x_i)^2$. Fix $\varepsilon = 0.001$ and consider two functions $f, \check{f} : \mathbb{H} \to \mathbb{R}$:

$$f(\boldsymbol{x}) = |\boldsymbol{x}|,$$

$$\check{f}(\boldsymbol{x}) = \max\left\{|\boldsymbol{x}|, \max_{n\geqslant 1}\{(1+\varepsilon)\cdot x_n - \tfrac{1}{n}\}\right\}.$$

Both of these functions are convex and Lipschitz, therefore their graphs in $\mathbb{H}\times\mathbb{R}$ equipped with its length metric form infinite-dimensional Alexandrov spaces, say $\mathcal{L}$ and $\check{\mathcal{L}}$ (this is proved formally in 13.7).

Let p be the origin of $\mathbb{H}\times\mathbb{R}$. Note that $\check{\mathcal{L}}\cap\mathcal{L}$ is starshaped in $\mathbb{H}$ with center at p. Further, $\check{\mathcal{L}}\backslash\mathcal{L}$ consists of a countable number of disjoint sets

$$\Omega_n = \left\{\, (\boldsymbol{x}, \check{f}(\boldsymbol{x})) \in \check{\mathcal{L}} : (1+\varepsilon)\cdot x_n - \tfrac{1}{n} > |\boldsymbol{x}| \,\right\}.$$

Note that $|\Omega_n - p| > \frac{1}{n}$ for each n. It follows that for any geodesic $[pq]$ in $\check{\mathcal{L}}$, a small subinterval $[p\bar{q}] \subset [pq]$ is a straight line segment in $\mathbb{H}\times\mathbb{R}$, and also a geodesic in $\mathcal{L}$. Thus we can treat $\Sigma_p\mathcal{L}$ and $\Sigma_p\check{\mathcal{L}}$ as one set, with two angle metrics $\measuredangle$ and $\check{\measuredangle}$. Let us denote by $\measuredangle_{\mathbb{H}\times\mathbb{R}}$ the angle in $\mathbb{H}\times\mathbb{R}$.

The space $\mathcal{L}$ is isometric to the Euclidean cone over $\Sigma_p\mathcal{L}$ with vertex at p; $\Sigma_p\mathcal{L}$ is isometric to a sphere in Hilbert space with radius $\frac{1}{\sqrt{2}}$. In particular, $\measuredangle$ is the length metric of $\measuredangle_{\mathbb{H}\times\mathbb{R}}$ on $\Sigma_p\mathcal{L}$.

Therefore in order to show that $\check{\measuredangle}$ does not define a length metric on $\Sigma_p\mathcal{L}$, it is sufficient to construct a pair of directions (ξ_+, ξ_-) such that

$$\check{\measuredangle}(\xi_+, \xi_-) < \measuredangle(\xi_+, \xi_-).$$

Set $\boldsymbol{e}_0 = (1, 0, 0, \ldots)$, $\boldsymbol{e}_1 = (0, 1, 0, \ldots), \ldots \in \mathbb{H}$. Consider the following two half-lines in $\mathbb{H}\times\mathbb{R}$:

$$\gamma_+(t) = \tfrac{t}{\sqrt{2}}\cdot(\boldsymbol{e}_0, 1) \quad\text{and}\quad \gamma_-(t) = \tfrac{t}{\sqrt{2}}\cdot(-\boldsymbol{e}_0, 1), \quad t\in[0, +\infty).$$

They form unit-speed geodesics in both $\mathcal{L}$ and $\check{\mathcal{L}}$. Let $\xi_\pm$ be the directions of $\gamma_\pm$ at p. Denote by σ_n the half-planes in $\mathbb{H}$ spanned by $\boldsymbol{e}_0$ and $\boldsymbol{e}_n$; that is, $\sigma_n = \{\, x\cdot\boldsymbol{e}_0 + y\cdot\boldsymbol{e}_n : y \geqslant 0\,\}$. Consider a sequence of 2-dimensional sectors $Q_n = \check{\mathcal{L}}\cap(\sigma_n\times\mathbb{R})$. For each n, the sector Q_n intersects Ω_n and is bounded by two geodesic half-lines $\gamma_\pm$. Note that $Q_n \xrightarrow[\text{GH}]{} Q$, where Q

is a solid Euclidean angle in $\mathbb{E}^2$ with angle measure $\beta < \measuredangle(\xi_+, \xi_-) = \frac{\pi}{\sqrt{2}}$. Indeed, Q_n is path-isometric to the subset of $\mathbb{E}^3$ described by

$$y \geqslant 0 \quad \text{and} \quad z = \max\left\{\sqrt{x^2+y^2}, (1+\varepsilon)\cdot y - \tfrac{1}{n}\right\}$$

with length metric. Thus its limit Q is path-isometric to the subset of $\mathbb{E}^3$ described by

$$y \geqslant 0 \quad \text{and} \quad z = \max\left\{\sqrt{x^2+y^2}, (1+\varepsilon)\cdot y\right\}$$

with length metric. In particular, for any $t, \tau \geqslant 0$,

$$\begin{aligned}|\gamma_+(t) - \gamma_-(\tau)|_{\check{\mathcal{L}}} \leqslant \lim_{n\to\infty} |\gamma_+(t) - \gamma_-(\tau)|_{Q_n} = \\ = \tilde{\gamma}^0\{\beta; t, \tau\}.\end{aligned}$$

That is, $\check{\measuredangle}(\xi_+, \xi_-) \leqslant \beta < \measuredangle(\xi_+, \xi_-)$. □

13.7. Lemma. *Let $\mathbb{H}$ be a Hilbert space, $f\colon \mathbb{H} \to \mathbb{R}$ be a convex Lipschitz function and $S \subset \mathbb{H}\times\mathbb{R}$ be the graph of f equipped with the length metric. Then S is* CBB(0)*.*

Proof. Recall that for a subset $X \subset \mathbb{H}\times\mathbb{R}$, we will denote by $|{*}-{*}|_X$ the length metric on X.

By the theorem of Buyalo [36], sharpened by the authors in [7], any convex hypersurface in a Euclidean space, equipped with the length metric, is non-negatively curved. Thus it is sufficient to show that for any 4-point set $\{x_0, x_1, x_2, x_3\} \subset S$, there is a finite-dimensional subspace $E \subset \mathbb{H}\times\mathbb{R}$ such that $\{x_i\} \in E$ and $|x_i - x_j|_{S\cap E}$ is arbitrary close to $|x_i - x_j|_S$.

Clearly $|x_i - x_j|_{S\cap E} \geqslant |x_i - x_j|_S$; thus it is sufficient to show that for given $\varepsilon > 0$ one can choose E so that

❷. $$|x_i - x_j|_{S\cap E} < |x_i - x_j|_S + \varepsilon.$$

For each pair (x_i, x_j), choose a broken line β_{ij} connecting x_i, x_j that lies under S (that is, outside of $\operatorname{Conv} S$) in $\mathbb{H}\times\mathbb{R}$ and has length at most $|x_i - x_j|_S + \varepsilon$. Let E be the affine hull of all the vertexes in all β_{ij}. Thus

$$|x_i - x_j|_{S\cap E} \leqslant \operatorname{length} \beta_{ij}$$

and ❷ follows. □

13.8. Exercise. *Construct a non-compact complete geodesic* CBB(0) *space that contains no half-lines.*

C Rademacher theorem

At the end of this section we give an extension of the Rademacher theorem (see Section 5D) to CBB and CAT spaces (13.12), proved by Alexander Lytchak [89]. The following proposition is the 1-dimensional case of the extended Rademacher theorem.

Recall that differentiable curves are defined in 6.10.

13.9. Proposition. *Let $\alpha\colon \mathbb{I}\to\mathcal{Z}$ be a locally Lipschitz curve in a complete length space. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT. *Then α is differentiable almost everywhere.*

The following two lemmas provide sufficient conditions for existence of the one-sided derivative of a curve in CBB and CAT spaces. The proofs of both lemmas are very similar.

13.10. Lemma. *Let $\alpha\colon \mathbb{I}\to\mathcal{L}$ be a 1-Lipschitz curve in a* CBB *space. Suppose that for some $t_0\in\mathbb{I}$ and any $\varepsilon>0$, there is a point p such that $|\alpha(t_0)-p|<\varepsilon$ and*

$$\varliminf_{t\to t_0+}\frac{\operatorname{dist}_p\circ\alpha(t)-\operatorname{dist}_p\circ\alpha(t_0)}{t-t_0}>1-\varepsilon.$$

Then the right derivative $\alpha^+(t_0)$ is defined and $|\alpha^+(t_0)|=1$.

Proof. Without loss of generality, we may assume that $t_0=0$. Set $x=$ $=\alpha(0)$. Fix a sequence of points $p_n\to x$ such that

$$\varliminf_{t\to 0+}\frac{|p_n-\alpha(t)|-|p_n-x|}{t}\to 1$$

as $n\to\infty$.

Observe that there are sequences $\delta_n\to 0+$ and $t_n\to 0+$ such that

❶ $$\tilde\measuredangle^\kappa(x\,{}^{\alpha(s)}_{p_n})>\pi-\delta_n\quad\text{and}\quad(1-\delta_n)\cdot s<|\alpha(s)-x|\leqslant s$$

for any $s\in(0,t_n]$.

For each n, choose $q_n\in\operatorname{Str}(x)$ sufficiently close $\alpha(t_n)$ that the inequality

$$\tilde\measuredangle^\kappa(x\,{}^{q_n}_{p_n})>\pi-\delta_n$$

still holds (see Definition 8.10).

Set $\gamma_n=\operatorname{geod}_{[xq_n]}$.

By comparison,

$$\begin{aligned}\tilde\measuredangle^\kappa(x\,{}^{\alpha(s)}_{\gamma_n(s)})&\leqslant 2\cdot\pi-\tilde\measuredangle^\kappa(x\,{}^{p_n}_{\gamma_n(s)})-\tilde\measuredangle^\kappa(x\,{}^{\alpha(s)}_{p_n})\leqslant\\&\leqslant 2\cdot\pi-\tilde\measuredangle^\kappa(x\,{}^{q_n}_{p_n})-\tilde\measuredangle^\kappa(x\,{}^{\alpha(s)}_{p_n})<\\&<2\cdot\delta_n.\end{aligned}$$

Therefore ❶ implies that

$$|\gamma_n(s) - \alpha(s)| < 10{\cdot}\delta_n{\cdot}(s)$$

if s is a sufficiently small and positive. That is, $\alpha^+(0)$ is defined (see Definition 6.7). □

13.11. Lemma. *Let $\alpha\colon \mathbb{I} \to \mathcal{U}$ be a 1-Lipschitz curve in a* CAT *space. Suppose that for some $t_0 \in \mathbb{I}$ and any $\varepsilon > 0$ there is a point q such that $|\alpha(t_0) - q| < \varepsilon$ and*

$$\varlimsup_{t\to t_0+} \frac{\operatorname{dist}_q\circ\alpha(t) - \operatorname{dist}_q\circ\alpha(t_0)}{t - t_0} < -1 + \varepsilon.$$

Then the right derivative $\alpha^+(t_0)$ is defined and $|\alpha^+(t_0)| = 1$.

Proof. Without loss of generatiy we may assume that $t_0 = 0$. Set $x = \alpha(0)$. Fix a sequence of points $q_n \to x$ such that

$$\varliminf_{t\to 0+} \frac{|q_n - \alpha(t)| - |q_n - x|}{t} \to -1$$

as $n \to \infty$.

Observe that there are sequences $\delta_n \to 0+$ and $t_n \to 0+$ such that

❷ $$\tilde\measuredangle^\kappa(x\,{}^{\alpha(s)}_{q_n}) < \delta_n \quad\text{and}\quad (1 - \delta_n){\cdot}s < |\alpha(s) - x| \leqslant s$$

for any $s \in (0, t_n]$.

Without loss of generality, we may assume that $|x - q_n| < \varpi^\kappa$ for any n; in particular, the geodesic $\gamma_n = \operatorname{geod}_{[xq_n]}$ is uniquely defined.

By comparison,

$$\begin{aligned}\tilde\measuredangle^\kappa(x\,{}^{\alpha(s)}_{\gamma_n(s)}) \leqslant \tilde\measuredangle^\kappa(x\,{}^{\alpha(s)}_{q_n}) <\\ < \delta_n.\end{aligned}$$

Therefore ❷ implies that

$$|\gamma_n(s) - \alpha(s)| < 10{\cdot}\delta_n{\cdot}s$$

if s is a sufficiently small and positive. That is, $\alpha^+(0)$ is defined (see Definition 6.7). □

Proof of 13.9. By the standard Rademacher theorem, we may assume that α has an arc-length parametrization. In particular α is 1-Lipschitz.

Recall that by Theorem 5.10,

❸ $$\operatorname{speed}_s\alpha \overset{a.e.}{=\!=} 1.$$

Fix a countable dense set $T \subset \mathbb{I}$; given $t \in T$, let

$$h_t(s) = |\alpha(t) - \alpha(s)|.$$

Note that h_t is 1-Lipschitz for each $t \in T$. Therefore, by the standard Rademacher theorem and countability of T for almost all $s \in \mathbb{I}$, $h_t'(s)$ is defined for all $t \in T$.

Let

$$w^+(s) := \varlimsup_{\substack{t \in T \\ t \to s-}} \{h_t'(s)\}.$$

Let us show that

❹ $$w^+(s) \stackrel{a.e.}{=\!=} 1.$$

Note that once this is proved, Lemma 13.10 implies the proposition in the CBB case.

For a small $\varepsilon > 0$, denote by N_ε^+ the set of all points $s \in \mathbb{I}$ such that $w^+(s) < 1 - \varepsilon$. Note that the sets N_ε^+ are measurable.

Suppose N_ε^+ has positive measure. Let $s_0 \in N_\varepsilon^+$ be a Lebesgue point of α. We may assume that $\operatorname{speed}_{s_0} \alpha = 1$ and $h_t'(s_0)$ is defined for any $t \in T$. Suppose $t \in T$ is sufficiently close to s_0 and $t < s_0$. Since $\operatorname{speed}_{s_0} \alpha = 1$, we have

❺ $$h_t(s_0) \geqslant (s_0 - t)\cdot(1 - \varepsilon^2).$$

Further, there is a set $A \subset [t, s_0]$ with measure at least $(1 - \varepsilon)\cdot|s_0 - t|$ such that

$$h_t'(s) < 1 - \varepsilon$$

for any $s \in A_n$. Since h_t is 1-Lipschitz, we have

$$\begin{aligned} h_t(s_0) &= \int\limits_{[t,s_0]\setminus A} h_t'(s)\cdot \boldsymbol{d}s + \int\limits_A h_t'(s)\cdot \boldsymbol{d}s \leqslant \\ &\leqslant (s_0 - t)\cdot[\varepsilon + (1 - \varepsilon)^2]. \end{aligned}$$

The latter contradicts ❺. Thus $w^+(s) \geqslant 1 - \varepsilon$ almost everywhere. Since $\varepsilon > 0$ is arbitrary, ❹ follows.

In the same way we can show that

❻ $$w^-(s) \stackrel{a.e.}{=\!=} -1,$$

where

$$w^-(s) := \varliminf_{\substack{t \in T \\ t \to s+}} \{h_t'(s)\}.$$

Then Lemma 13.11 implies the proposition in the CAT case. □

13.12. Extended Rademacher theorem. *Let $f\colon \mathbb{E}^m \looparrowright \mathcal{Z}$ be a locally Lipschitz submap from a Euclidean space to a complete length space $\mathcal{Z}$. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT. *Then the differential $\boldsymbol{d}_x f$ is defined at almost all points $x \in \operatorname{Dom} f$.*

Moreover the differential $\boldsymbol{d}_x f$ is linear at almost all x in the following sense: the image $\operatorname{Im} f$ is a convex subcone of $\mathrm{T}_{f(x)}\mathcal{Z}$, and there is an isometry ι from $\operatorname{Im} f$ to a Euclidean space such that the composition $\iota \circ \circ \boldsymbol{d}_x f$ is linear.

The proof is a reduction to the 1-dimensional case (13.9) by standard arguments [76, 92].

Proof. Without loss of generality, we may assume that $\operatorname{Dom} f$ is bounded and f is Lipschitz.

Fix a countable dense set of vectors $\{v_i\}$ in $\mathbb{E}^m$. Fix v_i and a point $p \in \in \operatorname{Dom} f$. By Proposition 13.9, the value $\boldsymbol{d}_x f(v_i)$ is defined at $x = p + t\cdot v_i$ for almost all t such that $x \in \operatorname{Dom} f$. It follows that $\boldsymbol{d}_x f(v_i)$ is defined for every i on a set A of full measure in $\operatorname{Dom} f$. Since the metric differential of f is defined almost everywhere (5.11), we have that $\boldsymbol{d}_x f(v)$ is defined for any v on a set B of full measure in $\operatorname{Dom} f$.

Applying the definitions of metric differential and differential (5.11 and 6.13), we obtain that the image of $\boldsymbol{d}_x f$ is a weakly convex set in $\mathrm{T}_{f(x)}$. It follows that $\operatorname{Im} \boldsymbol{d}_x f$ is CBB(0) or CAT(0) if the space $\mathcal{Z}$ is CBB or CAT respectively. It remains to apply Exercise 8.15 or 9.16 if the space $\mathcal{Z}$ is CBB or CAT respectively. □

D Differential

13.13. Exercise. *Let $\mathcal{U}$ be a complete length* CAT(κ) *space and $p, q \in \mathcal{U}$. Assume $|p - q| < \varpi^\kappa$. Show that*

$$(\boldsymbol{d}_q \operatorname{dist}_p)(v) = -\langle \uparrow_{[qp]}, v\rangle.$$

13.14. Exercise. *Let $\mathcal{L}$ be a length* CBB(κ) *space and $p, q \in \mathcal{L}$ be distinct points. Assume $q \in \operatorname{Str}(p)$ or $p \in \operatorname{Str}(q)$ Show that*

$$(\boldsymbol{d}_q \operatorname{dist}_p)(v) = -\langle \uparrow_{[qp]}, v\rangle.$$

13.15. Lemma. *Let $\mathcal{U}$ be a complete length* CAT *space, $f\colon \mathcal{U} \looparrowright \mathbb{R}$ be a locally Lipschitz semiconcave subfunction, and $p \in \operatorname{Dom} f$. Then $\boldsymbol{d}_p f$ is a Lipschitz concave function on $\mathrm{T}_p\mathcal{U}$.*

Proof. Recall that the tangent space $\mathrm{T}_p = T_p\mathcal{U}$ can be considered as a subspace of the ultratangent space T_p^{ω} (6.15). Since T_p^{ω} is CAT(0), 13.4 implies that T_p is a convex set in T_p^{ω}.

By 13.3, $\boldsymbol{d}_p^{\omega} f$ is a concave function on T_p^{ω}. It remains to apply that $\boldsymbol{d}_p f = (\boldsymbol{d}_p^{\omega} f)|_{\mathrm{T}_p}$ (6.14c). □

As is shown in Halbeisen's example (see Section 13B), a CBB space might have tangent spaces that are not length spaces; thus concavity of the differential $\boldsymbol{d}_p f$ of a semiconcave function f is meaningless. Nevertheless, as the following lemma shows, the differential $\boldsymbol{d}_p f$ of a semiconcave function always satisfies the following weaker property similar to concavity (compare [112, 136], [99, 4.2]). In the finite dimensional case, $\boldsymbol{d}_p f$ is concave.

13.16. Lemma. *Let $\mathcal{Z}$ be a complete length space, $f\colon \mathcal{Z} \hookrightarrow \mathbb{R}$ be a locally Lipschitz semiconcave subfunction, and $p \in \operatorname{Dom} f$. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT. *Then for any $u, v \in \mathrm{T}_p$, we have*

$$s\cdot\sqrt{|u|^2 + 2\cdot\langle u, v\rangle + |v|^2} \geqslant (\boldsymbol{d}_p f)(u) + (\boldsymbol{d}_p f)(v),$$

where

$$s = \sup\{\, (\boldsymbol{d}_p f)(\xi) \,:\, \xi \in \Sigma_p \,\}.$$

Proof of 13.16. If $\mathcal{Z}$ is CAT, then the statement follows from 13.15. Indeed, let z be the midpoint of a geodesic $[uv]_{\mathrm{T}_p}$. Observe that $2\cdot|z| = \sqrt{|u|^2 + 2\cdot\langle u, v\rangle + |v|^2}$. Since $\boldsymbol{d}_p f$ is concave, we have that

$$2\cdot\boldsymbol{d}_p f(z) \geqslant \boldsymbol{d}_p f(u) + \boldsymbol{d}_p f(v).$$

It remains to choose $\xi \in \Sigma_p$ so that $\xi\cdot|z| = z$ and observe that $s \geqslant \boldsymbol{d}_p(\xi)$.

Now assume $\mathcal{Z}$ is CBB. We can assume that $\alpha = \measuredangle(u, v) > 0$, otherwise the statement is trivial. Moreover, since $\mathrm{T}_p' = \operatorname{Cone}(\Sigma_p')$ is dense in T_p and $\boldsymbol{d}_p f\colon \mathrm{T}_p \to \mathbb{R}$ is Lipschitz, we can assume that $u, v \in \mathrm{T}_p'$; that is, $\exp_p(t\cdot u)$ and $\exp_p(t\cdot v)$ are defined for all small $t > 0$.

Prepare a model configuration of five points: $\tilde p$, $\tilde u$, $\tilde v$, $\tilde v'$, $\tilde w \in \mathbb{E}^2$ such that

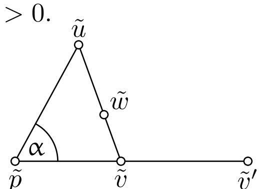

- ◇ $\measuredangle[\tilde p\,{}^{\tilde u}_{\tilde v}] = \alpha$,
- ◇ $|\tilde p - \tilde u| = |u|$,
- ◇ $|\tilde p - \tilde v| = |v|$,
- ◇ $\tilde v'$ lies on an extension of $[\tilde p\tilde v]$ so that $\tilde v$ is the midpoint of $[\tilde p\tilde v']$,
- ◇ $\tilde w$ is the midpoint between $\tilde u$ and $\tilde v$.

Note that

$$|\tilde p - \tilde w| = \tfrac{1}{2}\cdot\sqrt{|u|^2 + 2\cdot\langle u, v\rangle + |v|^2}.$$

Assume that $\mathcal{Z}$ is geodesic.

For all small $t > 0$, construct points $u_t, v_t, v_t', w_t \in \mathcal{Z}$ as follows:

(a) $v_t = \exp_p(t{\cdot}v)$, $\quad v'_t = \exp_p(t{\cdot}v')$
(b) $u_t = \exp_p(t{\cdot}u)$.
(c) w_t is the midpoint of $[u_t v_t]$.

Clearly $|p - u_t| = t{\cdot}|u|$, $|p - v_t| = t{\cdot}|v|$, $|p - v'_t| = t{\cdot}|v'|$. Since $\measuredangle(u,v)$ is defined, we have $|u_t - v_t| = t{\cdot}|\tilde u - \tilde v| + o(t)$ and $|u_t - v'_t| = t{\cdot}|\tilde u - \tilde v'| + o(t)$ (see Theorem 8.14c and Section 8A).

From the point-on-side and hinge comparisons (8.14b+c), we have

$$\tilde\measuredangle^{\kappa}(v_t\,{}^{p}_{w_t}) \geqslant \tilde\measuredangle^{\kappa}(v_t\,{}^{p}_{u_t}) \geqslant \measuredangle[\tilde v\,{}^{\tilde p}_{\tilde u}] + \tfrac{o(t)}{t}$$

and

$$\tilde\measuredangle^{\kappa}(v_t\,{}^{v'_t}_{w_t}) \geqslant \tilde\measuredangle^{\kappa}(v_t\,{}^{v'_t}_{u_t}) \geqslant \measuredangle[\tilde v\,{}^{\tilde v'}_{\tilde u}] + \tfrac{o(t)}{t}.$$

Clearly, $\measuredangle[\tilde v\,{}^{\tilde p}_{\tilde u}] + \measuredangle[\tilde v\,{}^{\tilde v'}_{\tilde u'}] = \pi$. From the adjacent angle comparison (8.14a), $\tilde\measuredangle^{\kappa}(v_t\,{}^{p}_{u_t}) + \tilde\measuredangle^{\kappa}(v_t\,{}^{u_t}_{v'_t}) \leqslant \pi$. Hence $\tilde\measuredangle^{\kappa}(v_t\,{}^{p}_{w_t}) \to \measuredangle[\tilde v\,{}^{\tilde p}_{\tilde w}]$ as $t \to 0+$ and thus

$$|p - w_t| = t|\tilde p - \tilde w| + o(t).$$

Since f is λ-concave we have

$$\begin{aligned} 2{\cdot}f(w_t) &\geqslant f(u_t) + f(v_t) + \tfrac{\lambda}{4}{\cdot}|u_t - v_t|^2 = \\ &= 2{\cdot}f(p) + t{\cdot}[(\boldsymbol{d}_p f)(u) + (\boldsymbol{d}_p f)(v)] + o(t). \end{aligned}$$

Applying λ-concavity of f, we have

❶ $$(\boldsymbol{d}_p f)(\uparrow_{[pw_t]}) \geqslant \frac{t{\cdot}[(\boldsymbol{d}_p f)(u) + (\boldsymbol{d}_p f)(v)] + o(t)}{2{\cdot}t{\cdot}|\tilde p - \tilde w| + o(t)}.$$

The lemma follows.

Finally, if $\mathcal{Z}$ is not geodesic one needs to make two adjustments in the above construction. Namely:

(i) For the geodesic $[u_t v_t]$ to be defined, in (b) one has to take $u_t \in \operatorname{Str}(v_t)$, $u_t \approx \exp_p(t{\cdot}u)$; more precisely,

$$|u_t - \exp_p(t{\cdot}u)| = o(t).$$

Thus instead of $|p - u_t| = t{\cdot}|u|$ we have

$$|p - u_t| = t|u| + o(t),$$

and this is sufficient for the rest of proof.

(ii) The direction $\uparrow_{[pw_t]}$ might be undefined. Thus in the estimate ❶, instead of $\uparrow_{[pw_t]}$ one should take $\uparrow_{[pw'_t]}$ for some point $w'_t \in \operatorname{Str}(p)$ near w_t (that is, $|w_t - w'_t| = o(t)$). □

E Gradient

13.17. Definition of gradient. *Let $\mathcal{X}$ be a length space with defined angles and $f\colon \mathcal{X} \hookrightarrow \mathbb{R}$ be a subfunction. Suppose for a point $p \in \operatorname{Dom} f$ the differential $\boldsymbol{d}_p f\colon \mathrm{T}_p \to \mathbb{R}$ is defined.*

A tangent vector $g \in \mathrm{T}_p$ is called a gradient of f at p (briefly, $g = \nabla_p f$) if

a) $(\boldsymbol{d}_p f)(w) \leqslant \langle g, w\rangle$ for any $w \in \mathrm{T}_p$, and

b) $(\boldsymbol{d}_p f)(g) = \langle g, g\rangle$.

13.18. Exercise. *Let $\mathcal{U}$ be a complete length* CAT(0) *space. Show that*

$$\nabla_p(-\operatorname{dist}_q) = \uparrow_{[pq]}$$

for any pair of distinct points $p, q \in \mathcal{U}$.

13.19. Existence and uniqueness of the gradient. *Let $\mathcal{Z}$ be a complete length space and $f\colon \mathcal{Z} \hookrightarrow \mathbb{R}$ be a locally Lipschitz and semiconcave subfunction. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT. *Then for any point $p \in \operatorname{Dom} f$, there is a unique gradient $\nabla_p f \in \mathrm{T}_p$.*

Proof; uniqueness. If $g, g' \in \mathrm{T}_p$ are two gradients of f, then

$$\langle g, g\rangle = (\boldsymbol{d}_p f)(g) \leqslant \langle g, g'\rangle, \qquad \langle g', g'\rangle = (\boldsymbol{d}_p f)(g') \leqslant \langle g, g'\rangle.$$

Therefore,

$$|g - g'|^2 = \langle g, g\rangle - 2\cdot\langle g, g'\rangle + \langle g', g'\rangle = 0,$$

that is, $g = g'$.

Existence. Note first that if $\boldsymbol{d}_p f \leqslant 0$, then one can take $\nabla_p f = 0$.

Otherwise, if $s = \sup\{\,(\boldsymbol{d}_p f)(\xi) : \xi \in \Sigma_p\,\} > 0$, it is sufficient to show that there is $\overline{\xi} \in \Sigma_p$ such that

❶ $$(\boldsymbol{d}_p f)\left(\overline{\xi}\right) = s.$$

Indeed, suppose $\overline{\xi}$ exists. Then applying Lemma 13.16 for $u = \overline{\xi}$, $v = \varepsilon\cdot w$ with $\varepsilon \to 0+$, we get

$$(\boldsymbol{d}_p f)(w) \leqslant \langle w, s\cdot\overline{\xi}\rangle$$

for any $w \in \mathrm{T}_p$; that is, $s\cdot\overline{\xi}$ is the gradient at p.

Take a sequence of directions $\xi_n \in \Sigma_p$, such that $(\boldsymbol{d}_p f)(\xi_n) \to s$. Applying Lemma 13.16 for $u = \xi_n$, $v = \xi_m$, we get

$$s \geqslant \frac{(\boldsymbol{d}_p f)(\xi_n) + (\boldsymbol{d}_p f)(\xi_m)}{\sqrt{2 + 2\cdot\cos\measuredangle(\xi_n, \xi_m)}}.$$

Therefore $\measuredangle(\xi_n, \xi_m) \to 0$ as $n, m \to \infty$; that is, (ξ_n) is Cauchy. Clearly $\overline{\xi} = \lim_n \xi_n$ satisfies ❶. □

Calculus

The next lemma states that the gradient points in the direction of maximal slope; moreover, if the slope in the given direction is almost maximal, then it is almost the direction of the gradient.

13.20. Lemma. *Let $\mathcal{Z}$ be a complete length space, $f\colon \mathcal{Z} \hookrightarrow \mathbb{R}$ be locally Lipschitz and semiconcave, and $p \in \operatorname{Dom} f$. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT.

Assume $|\nabla_p f| > 0$; let $\bar\xi = \frac{1}{|\nabla_p f|}\cdot\nabla_p f$. Then:

a) If for some $v \in \mathrm{T}_p$, we have

$$|v| \leqslant 1 + \varepsilon \quad \text{and} \quad (\boldsymbol{d}_p f)(v) > |\nabla_p f|\cdot(1-\varepsilon),$$

then

$$|\bar\xi - v| < 100\cdot\sqrt{\varepsilon}.$$

b) If $v_n \in \mathrm{T}_p$ is a sequence of vectors such that

$$\overline{\lim_{n\to\infty}}\, |v_n| \leqslant 1 \quad \text{and} \quad \underline{\lim_{n\to\infty}}\, (\boldsymbol{d}_p f)(v_n) \geqslant |\nabla_p f|,$$

then

$$\lim_{n\to\infty} v_n = \bar\xi.$$

c) $\bar\xi$ is the unique maximum direction for the restriction $\boldsymbol{d}_p f|_{\Sigma_p}$. In particular,

$$|\nabla_p f| = \sup\left\{\, \boldsymbol{d}_p f \;:\; \xi \in \Sigma_p f \,\right\}.$$

Proof. According to the definition of gradient,

$$\begin{aligned} |\nabla_p f|\cdot(1-\varepsilon) < (\boldsymbol{d}_p f)(v) &\leqslant \\ &\leqslant \langle v, \nabla_p f\rangle = \\ &= |v|\cdot|\nabla_p f|\cdot\cos\measuredangle(\nabla_p f, v). \end{aligned}$$

Thus $|v| > 1 - \varepsilon$ and $\cos\measuredangle(\nabla_p f, v) > \frac{1-\varepsilon}{1+\varepsilon}$. Hence (a).

Statements (b) and (c) follow directly from (a). □

As a corollary of the above lemma and Proposition 5.18 we obtain the following:

13.21. Chain rule. *Let $\mathcal{Z}$ be a complete length space, $f\colon \mathcal{Z} \hookrightarrow \mathbb{R}$ be a semiconcave function, and $\varphi\colon \mathbb{R} \to \mathbb{R}$ be a nondecreasing semiconcave function. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT. *Then $\varphi\circ f$ is semiconcave and $\nabla_x(\varphi\circ f) = \varphi^+(f(x))\cdot\nabla_x f$ for any $x \in \operatorname{Dom} f$.*

The following inequalities describe an important property of the "gradient vector field".

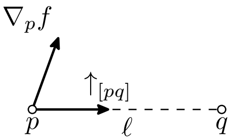

13.22. Lemma. *Let $\mathcal{Z}$ be a complete length space, $f\colon \mathcal{Z} \hookrightarrow \mathbb{R}$ satisfy $f'' + \kappa{\cdot}f \leqslant \lambda$ for some $\kappa, \lambda \in \mathbb{R}$. Let $[pq] \subset \operatorname{Dom} f$. and $\ell = |p-q|$. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT. *Then*

$$\langle \uparrow_{[pq]}, \nabla_p f\rangle \geqslant \frac{f(q) - f(p){\cdot}\mathrm{cs}^\kappa\ell - \lambda{\cdot}\mathrm{md}^\kappa\ell}{\mathrm{sn}^\kappa\ell}.$$

In particular,

a) if $\kappa = 0$,

$$\langle \uparrow_{[pq]}, \nabla_p f\rangle \geqslant \big(f(q) - f(p) - \tfrac{\lambda}{2}{\cdot}\ell^2\big)/\ell;$$

b) if $\kappa = 1$, $\lambda = 0$ we have

$$\langle \uparrow_{[pq]}, \nabla_p f\rangle \geqslant (f(q) - f(p){\cdot}\cos\ell)\,/\sin\ell;$$

c) if $\kappa = -1$, $\lambda = 0$ we have

$$\langle \uparrow_{[pq]}, \nabla_p f\rangle \geqslant (f(q) - f(p){\cdot}\cosh\ell)\,/\sinh\ell.$$

Proof of 13.22. Note that

$$\mathrm{geod}_{[pq]}(0) = p, \quad \mathrm{geod}_{[pq]}(\ell) = q, \quad (\mathrm{geod}_{[pq]})^+(0) = \uparrow_{[pq]}.$$

Thus,

$$\begin{aligned}\langle \uparrow_{[pq]}, \nabla_p f\rangle &\geqslant d_p f(\uparrow_{[pq]}) = \\ &= (f\circ \mathrm{geod}_{[pq]})^+(0) \geqslant \\ &\geqslant \frac{f(q) - f(p){\cdot}\mathrm{cs}^\kappa\ell - \lambda{\cdot}\mathrm{md}^\kappa\ell}{\mathrm{sn}^\kappa\ell}.\end{aligned}$$

□

The following corollary states that the gradient vector field is monotonic in a sense similar to the definition of monotone operators; see for example [108].

13.23. λ-Monotonicity of gradient. *Let $\mathcal{Z}$ be a complete length space, $f\colon \mathcal{Z} \hookrightarrow \mathbb{R}$ be locally Lipschitz and λ-concave and $[pq] \subset \operatorname{Dom} f$. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT. *Then*

$$\langle \uparrow_{[pq]}, \nabla_p f\rangle + \langle \uparrow_{[qp]}, \nabla_q f\rangle \geqslant -\lambda{\cdot}|p-q|.$$

Proof. Add two inequalities from 13.22*a*. □

13.24. Lemma. *Let $\mathcal{Z}$ be a complete length space, $f, g\colon \mathcal{Z} \looparrowright \mathbb{R}$, and $p \in \operatorname{Dom} f \cap \operatorname{Dom} g$. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT*.*

Then

$$|\nabla_p f - \nabla_p g|^2_{\mathrm{T}_p} \leqslant s\cdot(|\nabla_p f| + |\nabla_p g|),$$

where

$$s = \sup\{\,|(\boldsymbol{d}_p f)(\xi) - (\boldsymbol{d}_p g)(\xi)| : \xi \in \Sigma_p\,\}.$$

In particular, if $f_n\colon \mathcal{Z} \looparrowright \mathbb{R}$ is a sequence of locally Lipschitz and semiconcave subfunctions, $p \in \operatorname{Dom} f_n$ for each n, and $\boldsymbol{d}_p f_n$ converges uniformly on Σ_p, then the sequence $\nabla_p f_n \in \mathrm{T}_p$ converges.

Proof. Clearly for any $v \in \mathrm{T}_p$, we have

$$|(\boldsymbol{d}_p f)(v) - (\boldsymbol{d}_p g)(v)| \leqslant s\cdot|v|.$$

From the definition of gradient (13.17) we have:

$$\begin{aligned} (\boldsymbol{d}_p f)(\nabla_p g) &\leqslant \langle \nabla_p f, \nabla_p g\rangle, & (\boldsymbol{d}_p g)(\nabla_p f) &\leqslant \langle \nabla_p f, \nabla_p g\rangle,\\ (\boldsymbol{d}_p f)(\nabla_p f) &= \langle \nabla_p f, \nabla_p f\rangle, & (\boldsymbol{d}_p g)(\nabla_p g) &= \langle \nabla_p g, \nabla_p g\rangle. \end{aligned}$$

Therefore,

$$\begin{aligned} |\nabla_p f - \nabla_p g|^2 &= \langle \nabla_p f, \nabla_p f\rangle + \langle \nabla_p g, \nabla_p g\rangle - 2\cdot\langle \nabla_p f, \nabla_p g\rangle \leqslant \\ &\leqslant (\boldsymbol{d}_p f)(\nabla_p f) + (\boldsymbol{d}_p g)(\nabla_p g) - \\ &\quad - (\boldsymbol{d}_p f)(\nabla_p g) - (\boldsymbol{d}_p g)(\nabla_p f) \leqslant \\ &\leqslant s\cdot(|\nabla_p f| + |\nabla_p g|). \end{aligned}$$

□

13.25. Exercise. *Let $\mathcal{L}$ be an m-dimensional complete length* CBB(κ) *space, the function $f\colon \mathcal{L} \to \mathbb{R}$ be semiconcave and locally Lipschitz, and $\alpha\colon \mathbb{I} \to \mathcal{L}$ be a Lipschitz curve. Show that*

$$\langle \nabla_{\alpha(t)} f, \alpha^+(t)\rangle = (\boldsymbol{d}_{\alpha(t)} f)(\alpha^+(t))$$

for almost all $t \in \mathbb{I}$.

Semicontinuity

In this section we collect a few consequences of the following lemma.

13.26. Ultralimit of |gradient|. *Assume that*

- ⋄ *$(\mathcal{Z}_n)$ is a sequence of complete length spaces and $(\mathcal{Z}_n, p_n) \to \to (\mathcal{Z}_\omega, p_\omega)$ as $n \to \omega$. Suppose that all $\mathcal{Z}_n$ are either* CBB *or* CAT*.*

◇ $f_n\colon \mathcal{Z}_n \hookrightarrow \mathbb{R}$ *and* $f_\omega\colon \mathcal{Z}_\omega \hookrightarrow \mathbb{R}$ *are locally Lipschitz and* λ-*concave, and* $f_n \to f_\omega$ *as* $n \to \omega$.

◇ $x_n \in \operatorname{Dom} f_n$ *and* $x_n \to x_\omega \in \operatorname{Dom} f_\omega$ *as* $n \to \omega$.

Then

$$|\nabla_{x_\omega} f_\omega| \leqslant \lim_{n\to\omega} |\nabla_{x_n} f_n|.$$

Remarks. The inequality might be strict. For example, consider $\mathcal{Z}_n = \mathbb{R}$, $f_n(x) = -|x|$ and $x_n \to 0+$.

From the convergence of gradient curves (proved later in 16.16), one can deduce the following slightly stronger statement.

13.27. Proposition. *Assume that*

◇ $\mathcal{Z}_n$ *is a sequence of complete length spaces and* $(\mathcal{Z}_n, p_n) \to (\mathcal{Z}_\omega, p_\omega)$ *as* $n \to \omega$. *Suppose that all* $\mathcal{Z}_n$ *are either* CBB *or* CAT.

◇ $f_n\colon \mathcal{Z}_n \hookrightarrow \mathbb{R}$ *and* $f_\omega\colon \mathcal{Z}_\omega \hookrightarrow \mathbb{R}$ *are locally Lipschitz and* λ-*concave and* $f_n \to f_\omega$ *as* $n \to \omega$.

Then

$$|\nabla_{x_\omega} f_\omega| = \inf\{\lim_{n\to\omega} |\nabla_{x_n} f_n|\},$$

where infimum is taken for all sequences $x_n \in \operatorname{Dom} f_n$ *such that* $x_n \to x_\omega \in \operatorname{Dom} f_\omega$ *as* $n \to \omega$.

Proof of 13.26. Fix an $\varepsilon > 0$ and choose $y_\omega \in \operatorname{Dom} f_\omega$ sufficiently close to x_ω that

$$|\nabla_{x_\omega} f_\omega| - \varepsilon < \frac{f_\omega(y_\omega) - f_\omega(x_\omega)}{|x_\omega - y_\omega|}.$$

Choose $y_n \in \mathcal{Z}_n$ such that $y_n \to y_\omega$ as $n \to \omega$. Since $|x_\omega - y_\omega|$ is sufficiently small, the λ-concavity of f_n implies that

$$|\nabla_{x_\omega} f_\omega| - 2\cdot\varepsilon < (\boldsymbol{d}_{x_n} f_n)(\uparrow_{[x_n y_n]})$$

for ω-almost all n. Hence

$$|\nabla_{x_\omega} f_\omega| - 2\cdot\varepsilon \leqslant \lim_{n\to\omega} |\nabla_{x_n} f_n|.$$

Since $\varepsilon > 0$ is arbitrary, the proposition follows. □

Note that the distance-preserving map $\iota\colon \mathcal{Z} \hookrightarrow \mathcal{Z}^\omega$ induces an embedding

$$\boldsymbol{d}_p\iota\colon \mathrm{T}_p\mathcal{Z} \hookrightarrow \mathrm{T}_p\mathcal{Z}^\omega.$$

Thus, we can (and will) consider $\mathrm{T}_p\mathcal{Z}$ as a subcone of $\mathrm{T}_p\mathcal{Z}^\omega$.

13.28. Corollary. *Let $\mathcal{Z}$ be a complete length space and $f\colon \mathcal{Z} \hookrightarrow \mathbb{R}$ be a locally Lipschitz semiconcave subfunction. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT. *Then*

$$\nabla_x f = \nabla_x f^\omega.$$

for any point $x \in \operatorname{Dom} f$.

Proof. Note that

$$\begin{array}{ccc} \mathcal{Z} & \supset & \mathcal{Z}^\omega \\ \cup & & \cup \\ \operatorname{Dom} f & \subset & \operatorname{Dom} f^\omega \end{array}.$$

Applying 13.26 for $\mathcal{Z}_n = \mathcal{Z}$ and $x_n = x$, we get that $|\nabla_x f| \geqslant |\nabla_x f^\omega|$.

On the other hand, $f = f^\omega|_{\mathcal{Z}}$, hence $\boldsymbol{d}_p f = \boldsymbol{d}_p f^\omega|_{\mathrm{T}_p \mathcal{Z}}$. Thus from 13.20*c*, $|\nabla_x f| \leqslant |\nabla_x f^\omega|$. Therefore

$$❷ \qquad |\nabla_x f| = |\nabla_x f^\omega|$$

for any $x \in \mathcal{Z}$.

Further,

$$\begin{aligned} |\nabla_x f|^2 &= (\boldsymbol{d}_x f)(\nabla_x f) = \\ &= \boldsymbol{d}_x f^\omega(\nabla_x f) \leqslant \\ &\leqslant \langle \nabla_x f^\omega, \nabla_x f\rangle = \\ &= |\nabla_x f^\omega| \cdot |\nabla_x f| \cdot \cos\measuredangle(\nabla_x f^\omega, \nabla_x f). \end{aligned}$$

Together with ❷, this implies $\measuredangle(\nabla_x f^\omega, \nabla_x f) = 0$ and the statement follows. □

13.29. Semicontinuity of |gradient|. *Let $\mathcal{Z}$ be a complete length space and $f\colon \mathcal{Z} \hookrightarrow \mathbb{R}$ be a locally Lipschitz semiconcave subfunction. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT. *Then the function $x \mapsto |\nabla_x f|$ is lower-continuous; that is for any sequence $x_n \to x \in \operatorname{Dom} f$, we have*

$$|\nabla_x f| \leqslant \varliminf_{n\to\infty} |\nabla_{x_n} f|.$$

Proof. According to 13.28, $|\nabla_x f| = |\nabla_x f^\omega|$. Applying 13.26 for $x_n \to x$, we obtain

$$\lim_{n\to\omega} |\nabla_{x_n} f| \geqslant |\nabla_x f^\omega| = |\nabla_x f|.$$

Passing to an arbitrary subsequence of x_n gives the result. □

F Polar vectors

Here we give a corollary of Lemma 13.24. It will be used to prove basic properties of the tangent space.

13.30. Anti-sum lemma. *Let $\mathcal{L}$ be a complete length* CBB *space and $p\in\mathcal{L}$.*

Given two vectors $u,v\in \mathrm{T}_p$, there is a unique vector $w\in\mathrm{T}_p$ such that

$$\langle u,x\rangle+\langle v,x\rangle+\langle w,x\rangle\geqslant 0$$

for any $x\in\mathrm{T}_p$, and

$$\langle u,w\rangle+\langle v,w\rangle+\langle w,w\rangle=0.$$

If T_p were a length space, then the lemma would follow from the existence of the gradient (13.19), applied to the function $\mathrm{T}_p\to\mathbb{R}$ defined by $x\mapsto -(\langle u,x\rangle+\langle v,x\rangle)$. However, the tangent space T_p might be not a length space; see Halbeisen's example 13.6.

Applying the above lemma for $u=v$, we have the following statement.

13.31. Existence of polar vector. *Let $\mathcal{L}$ be a complete length* CBB *space and $p\in\mathcal{L}$. Given a vector $u\in\mathrm{T}_p$, there is a unique vector $u^*\in$ $\in\mathrm{T}_p$ such that $\langle u^*,u^*\rangle+\langle u,u^*\rangle=0$ and u^* is polar to u; that is, $\langle u^*,x\rangle+\langle u,x\rangle\geqslant 0$ for any $x\in\mathrm{T}_p$.*

In particular, for any vector $u\in\mathrm{T}_p$ there is a polar vector $u^*\in\mathrm{T}_p$ such that $|u^*|\leqslant|u|$.

Milka's lemma provides a refinement of this statement; it states that in the finite-dimensional case, we can assume that $|u^*|=|u|$.

It is instructive to solve the following exercise before reading the proof of 13.30.

13.32. Exercise. *Let $\mathcal{L}$ be a complete length* CBB(κ) *space and a,b,p be mutually distinct points in $\mathcal{L}$. Prove that*

$$(\boldsymbol{d}_p\operatorname{dist}_a)(\nabla_p\operatorname{dist}_b)\leqslant\cos\tilde\measuredangle^\kappa(p\,{}^a_b).$$

Proof of 13.30. Choose two sequences of points $a_n,b_n\in\operatorname{Str}(p)$ such that $\uparrow_{[pa_n]}\to u/|u|$ and $\uparrow_{[pb_n]}\to v/|v|$. Consider a sequence of functions

$$f_n=|u|\cdot\operatorname{dist}_{a_n}+|v|\cdot\operatorname{dist}_{b_n}.$$

According to Exercise 13.14,

$$(\boldsymbol{d}_pf_n)(x)=-|u|\cdot\langle\uparrow_{[pa_n]},x\rangle-|v|\cdot\langle\uparrow_{[pb_n]},x\rangle.$$

Thus we have the following uniform convergence for all $x \in \Sigma_p$:

$$(\boldsymbol{d}_p f_n)(x) \xrightarrow[n\to\infty]{} -\langle u, x\rangle - \langle v, x\rangle.$$

According to Lemma 13.24, the sequence $\nabla_p f_n$ converges. Let

$$w = \lim_n \nabla_p f_n.$$

By the definition of gradient,

$$\begin{aligned} \langle w, w\rangle &= \lim_{n\to\infty} \langle \nabla_p f_n, \nabla_p f_n\rangle = \\ &= \lim_{n\to\infty} (\boldsymbol{d}_p f_n)(\nabla_p f_n) = \\ &= -\langle u, w\rangle - \langle v, w\rangle, \end{aligned} \qquad \begin{aligned} \langle w, x\rangle &= \lim_{n\to\infty} \langle \nabla_p f_n, x\rangle \geqslant \\ &\geqslant \lim_{n\to\infty} (\boldsymbol{d}_p f_n)(x) = \\ &= -\langle u, x\rangle - \langle v, x\rangle. \end{aligned}$$

□

G Linear subspace of tangent space

13.33. Definition. *Let $\mathcal{L}$ be a complete length* CBB(κ) *space, $p \in \mathcal{L}$ and $u, v \in \mathrm{T}_p$. We say that vectors u and v are opposite to each other, (briefly, $u + v = 0$) if $|u| = |v| = 0$ or $\measuredangle(u, v) = \pi$ and $|u| = |v|$.*

The subcone

$$\mathrm{Lin}_p = \{\, v \in \mathrm{T}_p : \exists\ w \in \mathrm{T}_p \quad \text{such that} \quad w + v = 0 \,\}$$

will be called the linear subcone of T_p.

The reason for the name "linear" will become evident in Theorem 13.36.

13.34. Proposition. *Let $\mathcal{L}$ be a complete length* CBB *space and $p \in \mathcal{L}$. Given two vectors $u, v \in \mathrm{T}_p$, the following statements are equivalent:*

a) $u + v = 0$;

b) $\langle u, x\rangle + \langle v, x\rangle = 0$ for any $x \in \mathrm{T}_p$;

c) $\langle u, \xi\rangle + \langle v, \xi\rangle = 0$ for any $\xi \in \Sigma_p$.

Proof. The condition $u + v = 0$ is equivalent to

$$\langle u, u\rangle = -\langle u, v\rangle = \langle v, v\rangle;$$

thus $(b)\Rightarrow(a)$. Since T_p is isometric to a subset of T_p^ω, the splitting theorem (16.21) applied for T_p^ω gives $(a)\Rightarrow(b)$.

The equivalence $(b)\Leftrightarrow(c)$ is trivial. □

13.35. Proposition. *Let $\mathcal{L}$ be a complete length* CBB *space and $p \in \mathcal{L}$. Then for any three vectors $u, v, w \in \mathrm{T}_p$, if $u + v = 0$ and $u + w = 0$ then $v = w$.*

Proof. By Proposition 13.34, both v and w satisfy the condition in corollary 13.31. Hence the result. □

Let $u \in \mathrm{Lin}_p$; that is $u + v = 0$ for some $v \in \mathrm{T}_p$. Given $s < 0$, let

$$s\cdot u := (-s)\cdot v.$$

In this way we define multiplication of any vector in Lin_p by any real number (positive and negative). Proposition 13.35 implies that such multiplication is uniquely defined.

13.36. Theorem. *Let $\mathcal{L}$ be a complete length* CBB(κ) *space and $p \in \mathcal{L}$. Then* Lin_p *is a subcone of* T_p *isometric to a Hilbert space.*

Before proving the theorem, let us give a corollary.

13.37. Corollary. *Let $\mathcal{L}$ be a complete length* CBB(κ) *space and $p \in$ $\in \mathrm{Str}(x_1, x_2, \dots, x_n)$. Then there is a subcone $E \subset \mathrm{T}_p$ that is isometric to a Euclidean space such that $\log[px_i] \in E$ for every i.*

Proof. By the definition of Str (8.10), $\log[px_i] \in \mathrm{Lin}_p$ for each i. It remains to apply Theorem 13.36. □

The main difficulty in the proof of Theorem 13.36 comes from the fact that in general T_p is not a length space; see Habeisen's example (13.6). If the tangent space were a length space, the statement would follow directly from the splitting theorem (16.21). In fact the proof of Theorem 13.36 is very circuitous — we use the construction of the gradient, as well as the splitting theorem, namely its corollary (16.22). Thus in order to understand our proof one needs to read most of Chapter 16.

Proof of 13.36. First we show that Lin_p is a complete geodesic CBB(0) space.

Recall that T_p^ω is a complete geodesic CBB(0) space (see 3.6 and 13.1*a*) and Lin_p is a closed subset of T_p^ω. Thus, it is sufficient to show that the metric on Lin_p inherited from T_p^ω is a length metric.

Fix two vectors $x, y \in \mathrm{Lin}_p$. Let u and v be such that $u + \tfrac12\cdot x = 0$ and $v + \tfrac12\cdot y = 0$. Apply Lemma 13.30 to the vectors u and v; let $w \in \mathrm{T}_p$ denote the obtained tangent vector.

❶ *w is a midpoint of $[xy]$.*

Indeed, according to Lemma 13.30,

$$\begin{aligned}|w|^2 &= -\langle w, u\rangle - \langle w, v\rangle = \\ &= \tfrac12\cdot\langle w, x\rangle + \tfrac12\cdot\langle w, y\rangle.\end{aligned}$$

Therefore

$$\begin{aligned}|x-w|^2+|w-y|^2&=2{\cdot}|w|^2+|x|^2+|y|^2-2{\cdot}\langle w,x\rangle-2{\cdot}\langle w,y\rangle=\\&=|x|^2+|y|^2-\langle w,x\rangle-\langle w,y\rangle\leqslant\\&\leqslant|x|^2+|y|^2+\langle u,x\rangle+\langle v,x\rangle+\langle u,y\rangle+\langle v,y\rangle=\\&=\tfrac12{\cdot}|x|^2+\tfrac12{\cdot}|y|^2-\langle x,y\rangle=\\&=\tfrac12{\cdot}|x-y|^2.\end{aligned}$$

Thus $|x-w|=|w-y|=\frac12{\cdot}|x-y|$ and ❶ follows. △

Note that for any $v\in\mathrm{Lin}_p$ there is a line ℓ that contains v and 0. Therefore by 16.22, Lin_p is isometric to a Hilbert space. □

H Comments

13.38. Open question. *Let $\mathcal{L}$ be a proper length* CBB(κ) *space. Is it true that for any $p\in\mathcal{L}$, the tangent space T_p is a length space?*

Chapter 14

Dimension of CAT spaces

In this chapter we discuss constructions introduced by Bruce Kleiner [77].

The material of this chapter is used mostly for CAT spaces, but the results in section 14A find some applications for finite-dimensional CBB spaces as well.

A The case of complete geodesic spaces

The following construction gives a $\Bbbk$-dimensional submanifold for a given "nondegenerate" array of $\Bbbk+1$ strongly convex functions.

14.1. Definition. *For two real arrays $\boldsymbol{v}$, $\boldsymbol{w}\in\mathbb{R}^{\Bbbk+1}$, $\boldsymbol{v}=(v^0,v^1,\dots,v^{\Bbbk})$ and $\boldsymbol{w}=(w^0,w^1,\dots,w^{\Bbbk})$, we will write $\boldsymbol{v}\succcurlyeq\boldsymbol{w}$ if $v^i\geqslant w^i$ for each i.*

Given a subset $Q\subset\mathbb{R}^{\Bbbk+1}$, denote by $\operatorname{Up}Q$ the smallest upper set containing Q, and by $\operatorname{Min}Q$ the set of minimal elements of Q with respect to $\succcurlyeq$; that is,

$$\begin{aligned}\operatorname{Up}Q&=\left\{\,\boldsymbol{v}\in\mathbb{R}^{\Bbbk+1}:\exists\,\boldsymbol{w}\in Q\text{ such that }\boldsymbol{v}\succcurlyeq\boldsymbol{w}\,\right\},\\ \operatorname{Min}Q&=\{\,\boldsymbol{v}\in Q:\text{if }\boldsymbol{v}\succcurlyeq\boldsymbol{w}\in Q\text{ then }\boldsymbol{w}=\boldsymbol{v}\,\}.\end{aligned}$$

14.2. Definition. *Let $\boldsymbol{f}=(f^0,f^1,\dots,f^{\Bbbk})\colon\mathcal{X}\to\mathbb{R}^{\Bbbk+1}$ be a function array on a metric space $\mathcal{X}$. The set*

$$\operatorname{Web}\boldsymbol{f}:=\boldsymbol{f}^{-1}\left[\operatorname{Min}\boldsymbol{f}(\mathcal{X})\right]\subset\mathcal{X}$$

will be called the web of $\boldsymbol{f}$.

Given an array $\boldsymbol{f}=(f^0,f^1,\dots,f^{\Bbbk})$, we denote by $\boldsymbol{f}^{-i}$ the subarray of $\boldsymbol{f}$ with f^i removed; that is,

$$\boldsymbol{f}^{-i}:=(f^0,\dots,f^{i-1},f^{i+1},\dots,f^{\Bbbk}).$$

Clearly $\operatorname{Web}\boldsymbol{f}^{-i}\subset\operatorname{Web}\boldsymbol{f}$. Define the inner web of $\boldsymbol{f}$ as

$$\operatorname{InWeb}\boldsymbol{f}=\operatorname{Web}\boldsymbol{f}\backslash\left(\bigcup_i\operatorname{Web}\boldsymbol{f}^{-i}\right).$$

We say that a function array is nondegenerate if $\operatorname{InWeb}\boldsymbol{f}\neq\varnothing$.

Example. If $\mathcal{X}$ is a geodesic space, then $\operatorname{Web}(\operatorname{dist}_x,\operatorname{dist}_y)$ is the union of all geodesics from x to y, and

$$\operatorname{InWeb}(\operatorname{dist}_x,\operatorname{dist}_y)=\operatorname{Web}(\operatorname{dist}_x,\operatorname{dist}_y)\backslash\{x,y\}.$$

Barycenters. Let us denote by $\Delta^{\Bbbk}\subset\mathbb{R}^{\Bbbk+1}$ the standard $\Bbbk$-simplex; that is, $\boldsymbol{x}=(x^0,x^1,\dots,x^{\Bbbk})\in\Delta^{\Bbbk}$ if $\sum_{i=0}^{\Bbbk}x^i=1$ and $x^i\geqslant0$ for all i.

Let $\mathcal{X}$ be a metric space and $\boldsymbol{f}=(f^0,f^1,\dots,f^{\Bbbk})\colon\mathcal{X}\to\mathbb{R}^{\Bbbk+1}$ be a function array. Consider the map $\mathfrak{S}_{\boldsymbol{f}}\colon\Delta^{\Bbbk}\to\mathcal{X}$ defined by

$$\mathfrak{S}_{\boldsymbol{f}}(\boldsymbol{x})=\operatorname{MinPoint}\sum_{i=0}^{\Bbbk}x^i\cdot f^i,$$

where $\operatorname{MinPoint}f$ denotes a point of minimum of f. The map $\mathfrak{S}_{\boldsymbol{f}}$ will be called a barycentric simplex of $\boldsymbol{f}$. Note that for a general function array $\boldsymbol{f}$, the value $\mathfrak{S}_{\boldsymbol{f}}(\boldsymbol{x})$ might be undefined or nonuniquely defined.

It is clear from the definition that $\mathfrak{S}_{\boldsymbol{f}^{-i}}$ coincides with the restriction of $\mathfrak{S}_{\boldsymbol{f}}$ to the corresponding face of $\Delta^{\Bbbk}$.

14.3. Theorem. *Let $\mathcal{X}$ be a complete geodesic space and $\boldsymbol{f}=(f^0,f^1,\dots$ $\dots,f^{\Bbbk})\colon\mathcal{X}\to\mathbb{R}^{\Bbbk+1}$ be an array of strongly convex and locally Lipschitz functions. Then $\boldsymbol{f}$ defines a $C^{\frac12}$-embedding $\operatorname{Web}\boldsymbol{f}\hookrightarrow\mathbb{R}^{\Bbbk+1}$.*

Moreover,

a) *$W=\operatorname{Up}[\boldsymbol{f}(\mathcal{X})]$ is a convex closed subset of $\mathbb{R}^{\Bbbk+1}$, and $S=\partial_{\mathbb{R}^{\Bbbk+1}}W$ is a convex hypersurface in $\mathbb{R}^{\Bbbk+1}$.*

b)
$$\boldsymbol{f}(\operatorname{Web}\boldsymbol{f})=\operatorname{Min}W\subset S$$
and
$$\boldsymbol{f}(\operatorname{InWeb}\boldsymbol{f})=\operatorname{Int}_S(\operatorname{Min}W).$$

c) *The barycentric simplex $\mathfrak{S}_{\boldsymbol{f}}\colon\Delta^{\Bbbk}\to\mathcal{X}$ is a uniquely defined Lipshitz map and $\operatorname{Im}\mathfrak{S}_{\boldsymbol{f}}=\operatorname{Web}\boldsymbol{f}$. In particular $\operatorname{Web}\boldsymbol{f}$ is compact.*

d) *Let us equip $\Delta^{\Bbbk}$ with the metric induced by the ℓ^1-norm on $\mathbb{R}^{\Bbbk+1}$. Then the Lipschitz constant of $\mathfrak{S}_{\boldsymbol{f}}\colon\Delta^{\Bbbk}\to\mathcal{U}$ can be estimated in terms of positive lower bounds on $(f^i)''$ and Lipschitz constants of f^i in a neighborhood of $\operatorname{Web}\boldsymbol{f}$ for all i.*

In particular, by (a) and (b), InWeb $\boldsymbol{f}$ *is* $C^{\frac{1}{2}}$*-homeomorphic to an open set of* $\mathbb{R}^{\Bbbk}$*.*

The proof is preceded by a few preliminary statements.

14.4. Lemma. *Suppose* $\mathcal{X}$ *is a complete geodesic space and* $f\colon \mathcal{X}\to\mathbb{R}$ *is a locally Lipschitz, strongly convex function. Then the minimum point of* f *is uniquely defined.*

Proof. Without loss of generality, we can assume that f is 1-convex. In particular, the following claim holds:

❶ *if* z *is a midpoint of the geodesic* $[xy]$*, then*

$$s\leqslant f(z)\leqslant \tfrac12\cdot f(x)+\tfrac12\cdot f(y)-\tfrac18\cdot|x-y|^2,$$

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 $\ell\in\mathbb{R}$ be a Lipschitz constant of f in a neighborhood of p.

Consider the function $\varphi(t)=f\circ\mathrm{geod}_{[px]}(t)$. Clearly φ is 1-convex and $\varphi^+(0)\geqslant-\ell$. Setting $\ell=|p-x|$, we have

$$\begin{aligned}f(x)&=\varphi(\ell)\geqslant\\&\geqslant f(p)-\ell\cdot\ell+\tfrac12\cdot\ell^2\geqslant\\&\geqslant f(p)-\tfrac12\cdot\ell^2.\end{aligned}$$

In particular,

$$\begin{aligned}s&:=\inf\{\,f(x):x\in\mathcal{X}\,\}\geqslant\\&\geqslant f(p)-\tfrac12\cdot\ell^2.\end{aligned}$$

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 Cauchy. Thus p_n converges to a minimum point of f. □

14.5. Definition. *Let* Q *be a closed subset of* $\mathbb{R}^{\Bbbk+1}$*. A vector* $\boldsymbol{x}=(x^0,x^1,\dots,x^{\Bbbk})\in\mathbb{R}^{\Bbbk+1}$ *is* s u b n o r m a l *to* Q *at a point* $\boldsymbol{v}\in Q$ *if*

$$\langle\boldsymbol{x},\boldsymbol{w}-\boldsymbol{v}\rangle:=\sum_i x^i\cdot(w^i-v^i)\geqslant 0$$

for any $\boldsymbol{w} \in Q$*.*

14.6. Lemma. *Let* $\mathcal{X}$ *be a complete geodesic space and* $\boldsymbol{f} = (f^0, f^1, \dots$ $\dots, f^{\Bbbk})\colon \mathcal{X} \to \mathbb{R}^{\Bbbk+1}$ *be an array of strongly convex and locally Lipschitz functions. Let* $W = \operatorname{Up} \boldsymbol{f}(\mathcal{X})$*. Then:*

a) W *is a closed convex set, bounded below with respect to* $\succcurlyeq$*.*

b) If $\boldsymbol{x}$ *is a subnormal vector to* W*, then* $\boldsymbol{x} \succcurlyeq \boldsymbol{0}$*.*

c) $S = \partial_{\mathbb{R}^{\Bbbk+1}} W$ *is a complete convex hypersurface in* $\mathbb{R}^{\Bbbk+1}$*.*

Proof. Denote by $\bar{W}$ the closure of W.

Convexity of the f^i implies that for any two points $p, q \in \mathcal{X}$ and $t \in [0,1]$ we have

❷ $$(1-t)\cdot \boldsymbol{f}(p) + t\cdot \boldsymbol{f}(q) \succcurlyeq \boldsymbol{f} \circ \operatorname{path}_{[pq]}(t),$$

where $\operatorname{path}_{[pq]}$ denotes a geodesic path from p to q. Therefore W, as well as $\bar{W}$, are convex sets in $\mathbb{R}^{\Bbbk+1}$.

Let

$$w^i = \min\left\{\, f^i(x) : x \in \mathcal{X} \,\right\}.$$

By Lemma 14.4, w^i is finite for each i. Clearly $\boldsymbol{w} = (w^0, w^1, \dots, w^{\Bbbk})$ is a lower bound of $\bar{W}$ with respect to $\succcurlyeq$.

It is clear that W has nonempty interior, and $W \neq \mathbb{R}^{\Bbbk+1}$ since W is bounded below. Therefore $S = \partial_{\mathbb{R}^{\Bbbk+1}} W = \partial_{\mathbb{R}^{\Bbbk+1}} \bar{W}$ is a complete convex hypersurface in $\mathbb{R}^{\Bbbk+1}$.

Since $\bar{W}$ is closed and bounded below, we also have

❸ $$\bar{W} = \operatorname{Up}[\operatorname{Min} \bar{W}].$$

Choose an arbitrary $\boldsymbol{v} \in S$. Let $\boldsymbol{x} \in \mathbb{R}^{\Bbbk+1}$ be a subnormal vector to $\bar{W}$ at $\boldsymbol{v}$. In particular, $\langle \boldsymbol{x}, \boldsymbol{y}\rangle \geqslant 0$ for any $\boldsymbol{y} \succcurlyeq \boldsymbol{0}$; that is, $\boldsymbol{x} \succcurlyeq \boldsymbol{0}$.

Further, according to Lemma 14.4, the function $\sum_i x^i\cdot f^i$ has a uniquely defined minimum point, say p. Clearly

❹ $$\boldsymbol{v} \succcurlyeq \boldsymbol{f}(p) \quad \text{and} \quad \boldsymbol{f}(p) \in \operatorname{Min} W.$$

Note that for any $\boldsymbol{u} \in \bar{W}$ there is $\boldsymbol{v} \in S$ such that $\boldsymbol{u} \succcurlyeq \boldsymbol{v}$. Therefore ❹ implies

$$\bar{W} \subset \operatorname{Up}[\operatorname{Min} W] \subset W.$$

Hence $\bar{W} = W$; that is, W is closed. □

Proof of 14.3; (a)+(b). Without loss of generality, we may assume that all f^i are 1-convex.

Given $\boldsymbol{v} = (v^0, v^1, \dots, v^{\Bbbk}) \in \mathbb{R}^{\Bbbk+1}$, 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 14.4, $\Phi(\boldsymbol{v})$ is uniquely defined.

From the definition of web (14.2) we have $\Phi\circ\boldsymbol{f}(p) = p$ for any $p\in$ $\in\operatorname{Web}\boldsymbol{f}$; that is, Φ is a left inverse to the restriction $\boldsymbol{f}|_{\operatorname{Web}\boldsymbol{f}}$. In particular,

❺ $$\operatorname{Web}\boldsymbol{f} = \operatorname{Im}\Phi.$$

Given $\boldsymbol{v},\boldsymbol{w}\in\mathbb{R}^{\Bbbk+1}$, set $p=\Phi(\boldsymbol{v})$ and $q=\Phi(\boldsymbol{w})$. Since $h_{\boldsymbol{v}}$ and $h_{\boldsymbol{w}}$ are 1-convex, we have

$$h_{\boldsymbol{v}}(q)\geqslant h_{\boldsymbol{v}}(p)+\tfrac12\cdot|p-q|^2, \qquad h_{\boldsymbol{w}}(p)\geqslant h_{\boldsymbol{w}}(q)+\tfrac12\cdot|p-q|^2.$$

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 $C^{\frac12}$-continuous, or $\boldsymbol{f}|_{\operatorname{Web}\boldsymbol{f}}$ is a $C^{\frac12}$-embedding.

As in Lemma 14.6, let $W=\operatorname{Up}\boldsymbol{f}(\mathcal{X})$ and $S=\partial_{\mathbb{R}^{\Bbbk+1}}W$. Then S is a convex hypersurface in $\mathbb{R}^{\Bbbk+1}$. Clearly $\boldsymbol{f}(\operatorname{Web}\boldsymbol{f})=\operatorname{Min}W\subset S$. From the definition of inner web, we have $\boldsymbol{v}\in\boldsymbol{f}(\operatorname{InWeb}\boldsymbol{f})$ if and only if $\boldsymbol{v}\in$ $\in S$ and for any i there is $\boldsymbol{w}=(w^0,w^1,\dots,w^{\Bbbk})\in W$ such that $w^j<v^j$ for all $j\ne i$. Thus $\boldsymbol{f}(\operatorname{InWeb}\boldsymbol{f})$ is open in S. That is, $\operatorname{InWeb}\boldsymbol{f}$ is $C^{\frac12}$-homeomorphic to an open set in a convex hypersurface $S\subset\mathbb{R}^{\Bbbk+1}$, and hence to an open set of $\mathbb{R}^{\Bbbk}$, as claimed.

(c)+(d). Since f^i is 1-convex, for any $\boldsymbol{x}=(x^0,x^1,\dots,x^{\Bbbk})\in\Delta^{\Bbbk}$ the convex combination

$$\left(\sum_i x^i\cdot f^i\right):\mathcal{X}\to\mathbb{R}$$

is also 1-convex. Therefore, according to Lemma 14.4, the barycentric simplex $\mathfrak{S}_{\boldsymbol{f}}$ is uniquely defined on $\Delta^{\Bbbk}$.

For $\boldsymbol{x},\boldsymbol{y}\in\Delta^{\Bbbk}$, let

$$\begin{aligned} f_{\boldsymbol{x}} &= \sum_i x^i\cdot f^i, & f_{\boldsymbol{y}} &= \sum_i y^i\cdot f^i,\\ p &= \mathfrak{S}_{\boldsymbol{f}}(\boldsymbol{x}), & q &= \mathfrak{S}_{\boldsymbol{f}}(\boldsymbol{y}),\\ \ell &= |p-q|. \end{aligned}$$

Note the following:

◇ The function $\varphi(t) = f_{\boldsymbol{x}} \circ \mathrm{geod}_{[pq]}(t)$ has minimum at 0. Therefore $\varphi^+(0) \geqslant 0$

◇ The function $\psi(t) = f_{\boldsymbol{y}} \circ \mathrm{geod}_{[pq]}(t)$ has minimum at ℓ. Therefore $\psi^-(\ell) \geqslant 0$.

From 1-convexity of $f_{\boldsymbol{y}}$, we have $\psi^+(0) + \psi^-(\ell) + \ell \leqslant 0$.

Let ℓ be a Lipschitz constant for all f^i in a neighborhood $\Omega \ni p$. Then

$$\psi^+(0) \leqslant \varphi^+(0) + \ell{\cdot}\|\boldsymbol{x} - \boldsymbol{y}\|_1,$$

where $\|\boldsymbol{x} - \boldsymbol{y}\|_1 = \sum_{i=0}^{k} |x^i - y^i|$. That is, given $\boldsymbol{x} \in \Delta^{k}$, there is a constant ℓ such that

$$\begin{aligned}|\mathfrak{S}_{\boldsymbol{f}}(\boldsymbol{x}) - \mathfrak{S}_{\boldsymbol{f}}(\boldsymbol{y})| = \ell &\leqslant \\ &\leqslant \ell{\cdot}\|\boldsymbol{x} - \boldsymbol{y}\|_1\end{aligned}$$

for any $\boldsymbol{y} \in \Delta^{k}$. 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

$$|\mathfrak{S}_{\boldsymbol{f}}(\boldsymbol{y}) - \mathfrak{S}_{\boldsymbol{f}}(\boldsymbol{z})| = \ell \leqslant \ell{\cdot}\|\boldsymbol{y} - \boldsymbol{z}\|_1$$

for any $\boldsymbol{y}$ and $\boldsymbol{z}$ sufficiently close to $\boldsymbol{x}$; that is, $\mathfrak{S}_{\boldsymbol{f}}$ is locally Lipschitz. Since Δ^{k} is compact, $\mathfrak{S}_{\boldsymbol{f}}$ is Lipschitz.

Clearly $\mathfrak{S}_{\boldsymbol{f}}(\Delta^{k}) \subset \mathrm{Web}\,\boldsymbol{f}$. It remains to show that $\mathfrak{S}_{\boldsymbol{f}}(\Delta^{k}) \supset \mathrm{Web}\,\boldsymbol{f}$. According to Lemma 14.6, $W = \mathrm{Up}\,\boldsymbol{f}(\mathcal{X})$ is a closed convex set in $\mathbb{R}^{k+1}$. Let $p \in \mathrm{Web}\,\boldsymbol{f}$. Clearly $\boldsymbol{f}(p) \in \mathrm{Min}\,W \subset S = \partial_{\mathbb{R}^{k+1}} W$. Let $\boldsymbol{x}$ be a subnormal vector to W at $\boldsymbol{f}(p)$. According to Lemma 14.6, $\boldsymbol{x} \succcurlyeq \boldsymbol{0}$. Without loss of generality, we may assume that $\sum_i x^i = 1$; that is, $\boldsymbol{x} \in \Delta^{k}$. By Lemma 14.4, p is the unique minimum point of $\sum_i x^i{\cdot}f^i$; that is, $p = \mathfrak{S}_{\boldsymbol{f}}(\boldsymbol{x})$. □

B The case of CAT spaces

Let $\boldsymbol{a} = (a^0, a^1, \dots, a^{k})$ be a point array in a metric space $\mathcal{U}$. Recall that $\mathrm{dist}_{\boldsymbol{a}}$ denotes the distance map

$$(\mathrm{dist}_{a^0}, \mathrm{dist}_{a^1}, \dots, \mathrm{dist}_{a^{k}})\colon \mathcal{U} \to \mathbb{R}^{k+1},$$

which can be also regarded as a function array. The *radius of the point array* $\boldsymbol{a}$ is defined to be the radius of the set $\{a^0, a^1, \dots, a^{k}\}$; that is,

$$\mathrm{rad}\,\boldsymbol{a} = \inf\left\{\, r > 0 : \exists z \in \mathcal{U} \text{ such that } a^i \in \mathrm{B}(z, r) \text{ for any } i \,\right\}.$$

Fix $\kappa \in \mathbb{R}$. Let $\boldsymbol{a} = (a^0, a^1, \dots, a^{\bar{k}})$ be a point array of radius $< \frac{\varpi^\kappa}{2}$ in a metric space $\mathcal{U}$. Consider the function array $\boldsymbol{f} = (f^0, f^1, \dots, f^{\bar{k}})$ where

$$f^i(x) = \operatorname{md}^\kappa |a^i - x|.$$

Assuming the barycentric simplex $\mathfrak{S}_{\boldsymbol{f}}$ is defined, then $\mathfrak{S}_{\boldsymbol{f}}$ is called the κ-barycentric simplex for the point array $\boldsymbol{a}$; it will be denoted by $\mathfrak{S}^\kappa_{\boldsymbol{a}}$. The points $a^0, a^1, \dots, a^{\bar{k}}$ are called vertexes of the κ-barycentric simplex. Note that once we say the κ-barycentric simplex is defined, we automatically assume that $\operatorname{rad} \boldsymbol{a} < \frac{\varpi^\kappa}{2}$.

14.7. Theorem. *Let $\mathcal{U}$ be a complete length $\mathrm{CAT}(\kappa)$ space and $\boldsymbol{a} = (a^0, a^1, \dots a^{\bar{k}})$ be a point array with radius $< \frac{\varpi^\kappa}{2}$. Then:*

a) The κ-barycentric simplex $\mathfrak{S}^\kappa_{\boldsymbol{a}} \colon \Delta^{\bar{k}} \to \mathcal{U}$ is defined. Moreover, $\mathfrak{S}^\kappa_{\boldsymbol{a}}$ is a Lipschitz map, and if $\Delta^{\bar{k}}$ is equipped with the ℓ^1-metric, then its Lipschitz constant can be estimated in terms of κ and the radius of $\boldsymbol{a}$ (in particular it does not depend on $\bar{k}$).

b) $\mathrm{Web}(\operatorname{dist}_{\boldsymbol{a}}) = \operatorname{Im} \mathfrak{S}^\kappa_{\boldsymbol{a}}$. Moreover, if a closed convex set $K \subset \mathcal{U}$ contains all a^i, then $\mathrm{Web}(\operatorname{dist}_{\boldsymbol{a}}) \subset K$.

c) The restriction[1] $\operatorname{dist}_{\boldsymbol{a}-0} |_{\mathrm{InWeb}(\operatorname{dist}_{\boldsymbol{a}})}$ is an open $C^{\frac{1}{2}}$-embedding in $\mathbb{R}^{\bar{k}}$. Thus there is an inverse of $\operatorname{dist}_{\boldsymbol{a}-0} |_{\mathrm{InWeb}(\operatorname{dist}_{\boldsymbol{a}})}$, say $\Phi \colon \mathbb{R}^{\bar{k}} \hookrightarrow \mathcal{U}$.
The subfunction $f = \operatorname{dist}_{a^0} \circ \Phi$ is semiconvex and locally Lipschitz. Moreover, if $\kappa \leqslant 0$, then f is convex.

In particular, $\mathrm{Web}(\operatorname{dist}_{\boldsymbol{a}})$ is a compact set and $\mathrm{InWeb}(\operatorname{dist}_{\boldsymbol{a}})$ is $C^{\frac{1}{2}}$-homeomorphic to an open subset of $\mathbb{R}^{\bar{k}}$.

14.8. Definition. *The submap $\Phi \colon \mathbb{R}^{\bar{k}} \hookrightarrow \mathcal{X}$ of Theorem 14.7c will be called the $\operatorname{dist}_{\boldsymbol{a}}$-web embedding with brace $\operatorname{dist}_{a^0}$. The terminology invokes Theorem 14.7c.*

14.9. Definition. *Let $\mathcal{U}$ be a complete length $\mathrm{CAT}(\kappa)$ space and $\boldsymbol{a} = (a^0, a^1, \dots a^{\bar{k}})$ be a point array with radius $< \frac{\varpi^\kappa}{2}$. If $\mathrm{InWeb}(\operatorname{dist}_{\boldsymbol{a}})$ is nonempty, then the point array $\boldsymbol{a}$ is called nondegenerate.*

Lemma 14.11 will provide examples of nondegenerate point arrays, which can be used in Theorem 14.7*c*.

14.10. Corollary. *Let $\mathcal{U}$ be a complete length $\mathrm{CAT}(\kappa)$ space, $\boldsymbol{a} = (a^0, a^1, \dots a^m)$ be a nondegenerate point array of radius $< \frac{\varpi^\kappa}{2}$ in $\mathcal{U}$ and $\sigma = \mathfrak{S}^\kappa_{\boldsymbol{a}}$ be the corresponding κ-baricentric simplex. Then for some*

[1] Recall that $\operatorname{dist}_{\boldsymbol{a}-0}$ denotes the array $(\operatorname{dist}_{a^1}, \dots, \operatorname{dist}_{a^{\bar{k}}})$.

$\boldsymbol{x}\in\Delta^m$, the differential $\boldsymbol{d}_{\boldsymbol{x}}\sigma$ is linear and the image $\operatorname{Im}\boldsymbol{d}_{\boldsymbol{x}}\sigma$ forms a subcone isometric to an m-dimensional Euclidean space in the tangent cone $\mathrm{T}_{\sigma(\boldsymbol{x})}$.

Proof. Denote the distance map $\operatorname{dist}_{\boldsymbol{a}^{-0}}$ by $\tau\colon\mathcal{U}\to\mathbb{R}^m$.

According to Theorem 14.7, σ is Lipschitz and the distance map τ gives an open embedding of $\operatorname{InWeb}(\operatorname{dist}_{\boldsymbol{a}})=\sigma(\Delta^m)\backslash\sigma(\partial\Delta^m)$. Note that τ is Lipschitz. According to Rademacher's theorem (13.12), the differential $\boldsymbol{d}_{\boldsymbol{x}}(\tau\circ\sigma)$ is linear for almost all $\boldsymbol{x}\in\Delta^m$. Further, since $\operatorname{InWeb}(\operatorname{dist}_{\boldsymbol{a}})\neq\neq\varnothing$, the area formula implies that $\boldsymbol{d}_{\boldsymbol{x}}(\tau\circ\sigma)$ is surjective on a set of positive masure of points $\boldsymbol{x}\in\Delta^m$.

Note that $\boldsymbol{d}_{\boldsymbol{x}}(\tau\circ\sigma)=(\boldsymbol{d}_{\sigma(\boldsymbol{x})}\tau)\circ(\boldsymbol{d}_{\boldsymbol{x}}\sigma)$. Applying Rademacher's theorem again, we have linearity of $\boldsymbol{d}_{\boldsymbol{x}}\sigma$ for almost all $\boldsymbol{x}\in\Delta^m$; at these points $\operatorname{Im}\boldsymbol{d}_{\boldsymbol{x}}\sigma$ forms a subcone isometric to a Euclidean space in $\mathrm{T}_{\sigma(\boldsymbol{x})}$. Clearly the dimension of $\operatorname{Im}\boldsymbol{d}_{\boldsymbol{x}}(\tau\circ\sigma)$ is at least as big as the dimension of $\operatorname{Im}\boldsymbol{d}_{\boldsymbol{x}}\sigma$. Hence the result. □

Proof of 14.7. Fix $z\in\mathcal{U}$ and $r<\frac{\varpi^\kappa}{2}$ such that $|z-a^i|<r$ for all i. Note that the set $K\cap\overline{\mathrm{B}}[z,r]$ is convex, closed, and contains all a^i. Applying the theorem on short retract (Exercise 9.73), we get the second part of (b).

The remaining statements are proved first in the case $\kappa\leqslant 0$, and then the remaining case $\kappa>0$ is reduced to the case $\kappa=0$.

Case $\kappa\leqslant 0$. Consider the function array $f^i=\operatorname{md}^\kappa\circ\operatorname{dist}_{a^i}$. From the definition of web (14.2), it is clear that $\operatorname{Web}(\operatorname{dist}_{\boldsymbol{a}})=\operatorname{Web}\boldsymbol{f}$. Further, from the definition of κ-barycentric simplex, $\mathfrak{S}^\kappa_{\boldsymbol{a}}=\mathfrak{S}_{\boldsymbol{f}}$.

All the functions f^i are strongly convex (see 9.25b). Therefore (a), (b) and the first statements in (c) follow from Theorem 14.3.

Case $\kappa>0$. Applying rescaling, we may assume $\kappa=1$, so $\varpi^\kappa=\varpi^1=\pi$.

Let $\mathring{\mathcal{U}}=\operatorname{Cone}\mathcal{U}$. By 11.7$a$, $\mathring{\mathcal{U}}$ is CAT(0). Let us denote by ι the natural embedding of $\mathcal{U}$ as the unit sphere in $\mathring{\mathcal{U}}$, and by $\operatorname{proj}\colon\mathring{\mathcal{U}}\hookrightarrow\mathcal{U}$ the submap defined by $\operatorname{proj}(v)=\iota^{-1}(v/|v|)$ for all $v\neq 0$. Note that there is $z\in\mathcal{U}$ and $\varepsilon>0$ such that the set

$$K_\varepsilon=\left\{\,v\in\mathring{\mathcal{U}}:\langle\iota(z),v\rangle\geqslant\varepsilon\,\right\}$$

contains all $\iota(a^i)$. Then $0\notin K_\varepsilon$, and the set K_ε is closed and convex. The latter follows from Exercise 9.27, since $v\mapsto-\langle\iota(z),v\rangle$ is a Busemann function.

Denote by $\iota(\boldsymbol{a})$ the point array $(\iota(a^0),\iota(a^1),\dots,\iota(a^k))$ in $\mathring{\mathcal{U}}$. From the case $\kappa=0$, we get that $\operatorname{Im}\mathfrak{S}^0_{\iota(\boldsymbol{a})}\subset K_\varepsilon$. In particular $\operatorname{Im}\mathfrak{S}^0_{\iota(\boldsymbol{a})}\not\ni 0$ and thus $\operatorname{proj}\circ\mathfrak{S}^0_{\iota(\boldsymbol{a})}$ is defined. Direct calculations show

$$\mathfrak{S}^1_{\boldsymbol{a}}=\operatorname{proj}\circ\mathfrak{S}^0_{\iota(\boldsymbol{a})}\quad\text{and}\quad\operatorname{Web}(\operatorname{dist}_{\boldsymbol{a}})=\operatorname{proj}[\operatorname{Web}(\operatorname{dist}_{\iota(\boldsymbol{a})})].$$

Thus the case $\kappa = 1$ of the theorem is reduced to the case $\kappa = 0$, which is proved already. ☐

14.11. Lemma. *Let $\mathcal{U}$ be a complete length* CAT(κ) *space, $\boldsymbol{a} = (a^0, a^1, \dots a^{\bar{k}})$ be an array of radius $< \frac{\varpi^\kappa}{2}$, and $B^i = \overline{\mathrm{B}}[a^i, r^i]$ for some array of positive reals $(r^0, r^1, \dots, r^{\bar{k}})$. Assume that $\bigcap_i B^i = \varnothing$, but $\bigcap_{i \neq j} B^i \neq \varnothing$ for any j. Then $\boldsymbol{a}$ is nondegenerate.*

Proof. Without loss of generality, we may assume that $\mathcal{U}$ is geodesic and $\operatorname{diam} \mathcal{U} < \varpi^\kappa$. If not, choose $z \in \mathcal{U}$ and $r < \frac{\varpi^\kappa}{2}$ so that $|z - a^i| \leqslant r$ for each i, and consider $\overline{\mathrm{B}}[z, r]$ instead of $\mathcal{U}$. The latter can be done since $\overline{\mathrm{B}}[z, r]$ is convex and closed, so $\overline{\mathrm{B}}[z, r]$ is a complete length CAT(κ) space and $\mathrm{Web}(\mathrm{dist}_{\boldsymbol{a}}) \subset \overline{\mathrm{B}}[z, r]$; see 9.26 and 14.7*b*.

By Theorem 14.7, $\mathrm{Web}(\mathrm{dist}_{\boldsymbol{a}})$ is a compact set; therefore there is a point $p \in \mathrm{Web}(\mathrm{dist}_{\boldsymbol{a}})$ minimizing the function

$$f(x) = \max_i \{\mathrm{dist}_{B^i} x\} = \max\{0, |a^0 - x| - r^0, \dots, |a^{\bar{k}} - x| - r^{\bar{k}}\}.$$

By the definition of web (14.2), p is also the minimum point of f on $\mathcal{U}$. Let us prove the following claim:

❶ *$p \notin B^j$ for any j.*

Indeed, assume the contrary; that is,

❷ $$p \in B^j$$

for some j. Then p is a point of local minimum for the function

$$h^j(x) = \max_{i \neq j} \{\mathrm{dist}_{B^i} x\}.$$

Hence

$$\max_{i \neq j} \{\measuredangle[p\,{}^{x}_{a^i}]\} \geqslant \tfrac{\pi}{2}$$

for any $x \in \mathcal{U}$. From the angle comparison (9.14*c*), it follows that p is a global minimum of h^j and hence

$$p \in \bigcap_{i \neq j} B^i.$$

The latter and ❷ contradict $\bigcap_i B^i = \varnothing$. △

From the definition of web, it also follows that

$$\mathrm{Web}(\mathrm{dist}_{\boldsymbol{a}^{-j}}) \subset \bigcup_{i \neq j} B^i.$$

Indeed, if $q \in \bigcap_{i \neq j} B^i$ and $q' \notin \bigcup_{i \neq j} B^i$, then $|a_i - q| < |a_i - q'|$ for any $i \neq j$ therefore $q' \notin \mathrm{Web}(\mathrm{dist}_{\boldsymbol{a}^{-j}})$. Therefore the claim implies that $p \notin \mathrm{Web}(\mathrm{dist}_{\boldsymbol{a}^{-j}})$ for each j; that is, $p \in \mathrm{InWeb}(\mathrm{dist}_{\boldsymbol{a}})$. ☐

C Dimension

See Chapter 7 for definitions of various dimension-like invariants of metric spaces.

We start with two examples.

The first example shows that the dimension of complete length CAT spaces is not local; that is, such spaces might have open sets with different linear dimensions.

Such an example can be constructed by gluing at one point two Euclidean spaces of different dimensions. According to Reshetnyak's gluing theorem (9.38), this construction gives a CAT(0) space.

The second example provides a complete length CAT space with topological dimension 1 and arbitrary large Hausdorff dimension. Thus for complete length CAT spaces, one should not expect any relations between topological and Hausdorff dimensions except for the one provided by Szpilrajn's theorem (7.5).

To construct the second type of example, note that the completion of any metric tree has topological dimension 1 and is $\mathrm{CAT}(\kappa)$ for any κ. Start with a binary tree Γ, and a sequence $\varepsilon_n > 0$ such that $\sum_n \varepsilon_n < \infty$. Define the metric on Γ by prescribing the length of an edge from level n to level $n+1$ to be ε_n. For an appropriately chosen sequence ε_n, the completion of Γ will contain a Cantor set of arbitrarily large Hausdorff dimension.

The following is a version of a theorem proved by Bruce Kleiner [77], with an improvement made by Alexander Lytchak [89].

14.12. Theorem. *For any complete length* $\mathrm{CAT}(\kappa)$ *space* $\mathcal{U}$*, the following statements are equivalent:*

a) $\operatorname{LinDim}\mathcal{U} \geqslant m$.

b) For some $z \in \mathcal{U}$ *there is an array of* $m+1$ *balls* $B^i = \mathrm{B}(a^i, r^i)$ *with* $a^0, a^1, \dots, a^m \in \mathrm{B}(z, \frac{\varpi^\kappa}{2})$ *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 $C^{\frac{1}{2}}$*-embedding* $\Phi\colon \overline{\mathrm{B}}[1]_{\mathbb{E}^m} \hookrightarrow \mathcal{U}$*; that is,* Φ *is bi-Hölder with exponent* $\frac{1}{2}$*.*

d) There is a closed separable set $K \subset \mathcal{U}$ *such that*

$$\operatorname{TopDim} K \geqslant m.$$

Remarks. Theorem 14.15 gives a stronger version of part (*c*) in the finite-dimensional case. Namely, a complete length CAT space with linear dimension m admits a bi-Lipschitz embedding Φ of an open set of $\mathbb{R}^m$. Moreover, the Lipschitz constants of Φ can be made arbitrarily close to 1.

14.13. Corollary. *For any separable complete length* CAT *space $\mathcal{U}$, we have*

$$\operatorname{TopDim}\mathcal{U} = \operatorname{LinDim}\mathcal{U}.$$

Any simplicial complex can be equipped with a length metric such that each $\Bbbk$-simplex is isometric to the standard simplex

$$\Delta^{\Bbbk} = \left\{ (x_0, \dots, x_{\Bbbk}) \in \mathbb{R}^{\Bbbk+1} : x_i \geqslant 0, \quad x_0 + \ldots + x_{\Bbbk} = 1 \right\}$$

with the metric induced by the ℓ^1-norm on $\mathbb{R}^{\Bbbk+1}$. This metric will be called the ℓ^1-metric on the simplicial complex.

14.14. Lemma. *Let $\mathcal{U}$ be a complete length* CAT(κ) *space and $\rho\colon \mathcal{U} \to \mathbb{R}$ be a continuous positive function. Then there is a simplicial complex $\mathcal{N}$ equipped with ℓ^1-metric, 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 LinDim$\,\mathcal{U} \leqslant m$ *then the Ψ-image of any closed simplex in $\mathcal{N}$ coincides with the image of its m-skeleton.*

Proof. Without loss of generality, we may assume that for any x we have $\rho(x) < \rho_0$ for some fixed $\rho_0 < \frac{\varpi^\kappa}{2}$.

By Stone's theorem, any metric space is paracompact. Thus, we can choose a locally finite covering $\{ \Omega_\alpha : \alpha \in \mathcal{A} \}$ of $\mathcal{U}$ such that $\Omega_\alpha \subset \mathrm{B}(x, \frac{1}{3}\cdot\rho(x))$ 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 $\{\alpha^0, \alpha^1, \dots, \alpha^n\} \subset \mathcal{A}$ are vertexes of a simplex if and only if $\Omega_{\alpha^0} \cap \Omega_{\alpha^1} \cap \ldots \cap \Omega_{\alpha^n} \neq \varnothing$.

Fix 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 κ-barycentric on each simplex. According to Theorem 14.7*a*, this extension exists, Ψ is Lipschitz, and its Lipschitz constant depends only on ρ_0 and κ.

(a) Fix $x \in \mathcal{U}$. Denote by Δ the minimal simplex that contains $\Phi(x)$, and let $\alpha^0, \alpha^1, \dots, \alpha^n$ be the vertexes of Δ. Note that α is a vertex of Δ if and only if $\varphi_\alpha(x) > 0$. Thus

$$|x - x_{\alpha^i}| < \tfrac{1}{3}\cdot\rho(x)$$

for any i. Therefore

$$\operatorname{diam}\Psi(\Delta) \leqslant \max_{i,j}\{|x_{\alpha^i} - x_{\alpha^j}|\} < \tfrac{2}{3}\cdot\rho(x).$$

In particular,

$$|x - \Psi\circ\Phi(x)| \leqslant |x - x_{\alpha^0}| + \operatorname{diam}\Psi(\Delta) < \rho(x).$$

(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 $\hbar > m$, there is a $\hbar$-simplex $\Delta^\hbar$ in $\mathcal{N}$ such that the barycentric simplex $\sigma = \Psi|_{\Delta^\hbar}$ is nondegenerate; that is,

$$W = \Psi(\Delta^\hbar)\backslash\Psi(\partial\Delta^\hbar) \neq \varnothing.$$

Applying Corollary 14.10 gives $\operatorname{LinDim}\mathcal{U} \geqslant \hbar$, a contradiction. □

Proof of 14.12. Note that

- ◇ The implication $(b)\Rightarrow(c)$ follows directly from Lemma 14.11 and Theorem 14.7c.
- ◇ The implication $(c)\Rightarrow(d)$ is trivial.

(d)⇒(a). According to Theorem 7.7, there is a continuous map $f\colon K \to \mathbb{R}^m$ with a stable value. By the Tietze extension theorem, it is possible to extend f to a continuous map $F\colon \mathcal{U} \to \mathbb{R}^m$.

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 Lemma 14.14 for the function ρ. For the resulting simplicial complex $\mathcal{N}$ and 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 $x \in \mathcal{U}$.

According to Lemma 5.5, there is a locally Lipschitz map $F_\varepsilon\colon \mathcal{U} \to \mathbb{R}^{m+1}$ such that $|F_\varepsilon(x) - F(x)| < \tfrac{1}{3}\cdot\varepsilon$ for any $x \in \mathcal{U}$.

Note that $\Phi(K)$ is contained in a countable subcomplex of $\mathcal{N}$, say $\mathcal{N}'$. Indeed, since K is separable, there is a countable dense collection of points $\{x_n\}$ in K. Denote by Δ_n the minimal simplex of $\mathcal{N}$ that contains $\Phi(x_n)$. Then $\Phi(K) \subset \bigcup_i \Delta_n$.

Arguing by contradiction, assume $\operatorname{LinDim}\mathcal{U} < m$. By 14.14*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 the restriction $F_\varepsilon\circ\Psi\circ\Phi|_K$. Since $\varepsilon>0$ is arbitrary, then $\mathbf{0}\in\mathbb{R}^m$ is not a stable value of $f=F|_K$, a contradiction.

(a)⇒(b). The following claim is a consequence of the definition of tangent space.

❶ *Let $q\in\mathcal{U}$ and $\dot x^1,\dot x^2,\dots,\dot x^n\in\mathrm{T}_q$. Then given $\delta>0$, there is an array $(x^1,x^2,\dots,x^n)$ of points in $\mathcal{U}$ such that*

$$\measuredangle(\dot x^i,\log[qx^i])<\delta,$$

and for some fixed $\mathfrak{z}>0$ we have

$$\tfrac{1}{\mathfrak{z}}\cdot|q-x^i|=|\dot x^i|\quad\textit{and}\quad\left|\tfrac{1}{\mathfrak{z}}\cdot|x^i-x^j|-|\dot x^i-\dot x^j|\right|<\delta$$

for all i and j.

Moreover the value $\mathfrak{z}$ can be taken arbitrarily small.

Proof of the claim. For each i choose a geodesic γ^i from q that goes almost in the directions of $\dot x^i$. Then take the point x^i on γ^i at distance $\mathfrak{z}\cdot|\dot x^i|$ from q. △

Choose $q\in\mathcal{U}$ such that T_q contains a subcone E isometric to m-dimensional Euclidean space. Note that one can choose $\varepsilon>0$ and a point arrray $(\dot a^0,\dot a^1,\dots,\dot a^m)$ in $E\subset\mathrm{T}_q$ such that $\bigcap_i\overline{\mathrm{B}}[\dot a^i,1+\varepsilon]=\varnothing$ and $\bigcap_{i\neq j}\overline{\mathrm{B}}[\dot a^i,1-\varepsilon]\neq\varnothing$ for each j.

Applying Claim ❶, we get a point array $(a^0,a^1,\dots,a^m)$ in $\mathcal{U}$ such that $\bigcap_i\overline{\mathrm{B}}[a^i,\mathfrak{z}]=\varnothing$ and $\bigcap_{i\neq j}\overline{\mathrm{B}}[a^i,\mathfrak{z}]\neq\varnothing$ for each j. Since $\mathfrak{z}>0$ can be chosen arbitrarily small, (b) follows. □

D Finite-dimensional spaces

Recall that a web embedding and its brace are defined in 14.8.

14.15. Theorem. *Suppose $\mathcal{U}$ is a complete length $\mathrm{CAT}(\kappa)$ space such that $\operatorname{LinDim}\mathcal{U}=m$, and $\boldsymbol{a}=(a^0,a^1,\dots a^m)$ is a point array in $\mathcal{U}$ with radius $<\frac{\varpi^\kappa}{2}$. Then the $\operatorname{dist}_{\boldsymbol{a}}$-web embedding $\Phi\colon\mathbb{R}^m\hookrightarrow\mathcal{U}$ with brace $\operatorname{dist}_{a^0}$ is locally Lipschitz.*

Note that if $\boldsymbol{a}$ is degenerate, that is, if $\operatorname{InWeb}(\operatorname{dist}_{\boldsymbol{a}})=\varnothing$, then the domain of the web embedding Φ above is empty and hence the conclusion of the theorem trivially holds.

14.16. Lemma. *Let $\mathcal{U}$ be a complete length* CAT(κ) *space, and* $\boldsymbol{a} = (a^0, a^1, \dots a^{\Bbbk})$ *be a point array with radius* $< \frac{\varpi^\kappa}{2}$. *Then for any* $p \in \mathrm{InWeb}(\mathrm{dist}_{\boldsymbol{a}})$, *there is* $\varepsilon > 0$ *such that if for some* $q \in \mathrm{Web}(\mathrm{dist}_{\boldsymbol{a}})$ *and* $b \in \mathcal{U}$ *we have*

$$|p-q| < \varepsilon, \quad |p-b| < \varepsilon \quad \textit{and} \quad \measuredangle[q\,{}^{b}_{a^i}] < \tfrac{\pi}{2} + \varepsilon$$

for each i, *then the array* $(b, a^0, a^1, \dots, a^m)$ *is nondegenerate.*

Proof. Without loss of generality, we may assume that $\mathcal{U}$ is geodesic and $\operatorname{diam}\mathcal{U} < \varpi^\kappa$. If not, consider instead of $\mathcal{U}$, a ball $\overline{\mathrm{B}}[z,r] \subset \mathcal{U}$ for some $z \in \mathcal{U}$ and $r < \frac{\varpi^\kappa}{2}$ such that $|z - a^i| \leqslant r$ for each i.

From the angle comparison (9.14c), it follows that $p \in \mathrm{InWeb}\,\boldsymbol{a}$ if and only if both of the following conditions hold:

1. $\max_i\{\measuredangle[p\,{}^{a^i}_{u}]\} \geqslant \frac{\pi}{2}$ for any $u \in \mathcal{U}$,
2. for each i there is $u^i \in \mathcal{U}$ such that $\measuredangle[p\,{}^{a^j}_{u^i}] < \frac{\pi}{2}$ for all $j \neq i$.

Due to the semicontinuity of angles (9.33), there is $\varepsilon > 0$ such that for any $x \in \mathrm{B}(p, 10\cdot\varepsilon)$ we have

❶ $$\measuredangle[x\,{}^{a^j}_{u^i}] < \tfrac{\pi}{2} - 10\cdot\varepsilon \quad \text{for all} \quad j \neq i.$$

Now assume that for sufficiently small $\varepsilon > 0$ there are points $b \in \mathcal{U}$ and $q \in \mathrm{Web}(\mathrm{dist}_{\boldsymbol{a}})$ such that

❷ $$|p-q| < \varepsilon, \quad |p-b| < \varepsilon, \quad \measuredangle[q\,{}^{b}_{a^i}] < \tfrac{\pi}{2} + \varepsilon \quad \text{for all} \quad i.$$

According to Theorem 14.7b, for all small $\varepsilon > 0$ we have

$$\mathrm{rad}\{b, a^0, a^1, \dots, a^{\Bbbk}\} < \tfrac{\varpi^\kappa}{2}.$$

Fix a sufficiently small $\delta > 0$ and let

$$v^i = \mathrm{geod}_{[qu^i]}(\tfrac{1}{3}\cdot\delta) \quad \text{and} \quad w^i = \mathrm{geod}_{[v^i b]}(\tfrac{2}{3}\cdot\delta).$$

Clearly

$$\begin{aligned}|b - w^i| &= |b - v^i| - \tfrac{2}{3}\cdot\delta \leqslant \\ &\leqslant |b-q| - \tfrac{1}{3}\cdot\delta.\end{aligned}$$

Further, the inequalities ❶ and ❷ imply

$$\begin{aligned}|a^j - w^i| &< |a^j - v^i| + \tfrac{2}{3}\cdot\varepsilon\cdot\delta < \\ &< |a^i - q| - \varepsilon\cdot\delta < \\ &< |a^i - q|\end{aligned}$$

for all $i \neq j$.

Set $B^i = \overline{\mathrm{B}}[a^i, |a^i - q|]$ and $B^{m+1} = \overline{\mathrm{B}}[b, |a^i - q| - \frac{1}{3}\cdot\delta]$. Clearly

$$\bigcap_{i \neq m+1} B^i = \{q\},$$
$$\bigcap_{i \neq j} B^i \ni w^j \quad \text{for} \quad j \neq m+1,$$
$$\bigcap_{i} B^i = \{q\} \cap B^{m+1} = \varnothing.$$

Lemma 14.11 finishes the proof. □

Proof of 14.15. Suppose Φ is not locally Lipshitz.; that is, there are sequences $\boldsymbol{y}_n, \boldsymbol{z}_n \to \boldsymbol{x} \in \operatorname{Dom}\Phi$ such that

$$❸ \qquad \frac{|\Phi(\boldsymbol{y}_n) - \Phi(\boldsymbol{z}_n|)}{|\boldsymbol{y}_n - \boldsymbol{z}_n|} \to \infty \quad \text{as} \quad n \to \infty.$$

Set $p = \Phi(\boldsymbol{x})$, $q_n = \Phi(\boldsymbol{y}_n)$, and $b_n = \Phi(\boldsymbol{z}_n)$. By 14.8, p, q_n, $b_n \in$ $\in \operatorname{InWeb}(\operatorname{dist}_{\boldsymbol{a}})$ and $q_n, b_n \to p$ as $n \to \infty$. Choose an arbitrary $\varepsilon > 0$. Note that ❸ implies

$$\measuredangle[q_n\,{}^{a^i}_{b_n}] < \tfrac{\pi}{2} + \varepsilon$$

for all $i > 0$ and all large n. Further, according to 14.8, the subfunction $(\operatorname{dist}_{\boldsymbol{a}^0}) \circ \Phi$ is locally Lipschitz. Therefore we also have

$$\measuredangle[q_n\,{}^{a^0}_{b_n}] < \tfrac{\pi}{2} + \varepsilon$$

for all large n. According to Lemma 14.16, the point array $b_n, a^0, \dots, a^{\Bbbk}$ for large n is nondegenerate.

Applying Corollary 14.10, we have a contradiction. □

E Remarks and open problems

The following conjecture (in an equivalent form) appears in [77], see also [59, p. 133].

14.17. Conjecture. *For any complete length* CAT *space* $\mathcal{U}$*, we have*

$$\operatorname{TopDim}\mathcal{U} = \operatorname{LinDim}\mathcal{U}.$$

By Corollary 14.13, this conjecture holds for separable spaces.

Chapter 15

Dimension of CBB spaces

In sections 15C and 15E, we prove equivalence of some dimension-like invariants for CBB spaces. In Section 14C we discuss the dimension of CAT spaces.

As the main dimension-like invariant, we will use the linear dimension LinDim; see Definition 7.9. In other words, by default *dimension* means *linear dimension*.

A Struts and rank

Our definitions of strut and distance chart differ from the one given by Burago, Gromov and Perelman [34]; it is closer to Perelman's definitions [102, 103].

The term "strut" seems to have the closest meaning to the original Russian term used in [34]. In the official translation, it appears as "burst", and in the authors' translation it was "strainer". Neither seems intuitive, so we decided to switch to "strut".

15.1. Definition of struts. *Let $\mathcal{L}$ be a complete length* CBB *space. We say that a point array $(a^0, a^1, \dots, a^{\bar{k}})$ in $\mathcal{L}$ is κ-strutting for a point $p \in \in \mathcal{L}$ if $\tilde\measuredangle^\kappa(p\,{}^{a^i}_{a^j}) > \frac{\pi}{2}$ for all $i \ne j$.*

The following definition is motivated by the observation that $k = = \operatorname{pack}_{\pi/2}(\mathbb{S}^{k-1}) - 1$ for any integer $k > 0$; the packing number is defined in 2B.

15.2. Definition. *Let $\mathcal{L}$ be a complete length* CBB *space and $p \in \mathcal{L}$. Let us define rank of $\mathcal{L}$ at p as*

$$\operatorname{rank}_p = \operatorname{rank}_p \mathcal{L} := \operatorname{pack}_{\pi/2} \Sigma_p - 1.$$

Thus rank takes values in $\mathbb{Z}_{\geqslant 0} \cup \{\infty\}$.

15.3. Proposition. *Let $\mathcal{L}$ be a complete length* CBB(κ) *space and $p \in \mathcal{L}$. Then the following conditions are equivalent:*

a) $\operatorname{rank}_p \geqslant \Bbbk$,

b) there is a point array $(a^0, a^1, \dots, a^\Bbbk)$ that is κ-strutting at p.

Proof of 15.3, (b)⇒(a). For each i, choose a point $\acute{a}^i \in \operatorname{Str}(p)$ sufficiently close to a^i (so $[p\acute{a}^i]$ exists for each i). One can choose $\acute{a}^i$ so that we still have $\tilde{\measuredangle}^\kappa(p\,{}^{\acute{a}^i}_{\acute{a}^j}) > \frac{\pi}{2}$ for all $i \neq j$.

From hinge comparison (8.14c),

$$\measuredangle(\uparrow_{[p\acute{a}^j]}, \uparrow_{[p\acute{a}^j]}) \geqslant \tilde{\measuredangle}^\kappa(p\,{}^{\acute{a}^i}_{\acute{a}^j}) > \tfrac{\pi}{2}$$

for all $i \neq j$. In particular $\operatorname{pack}_{\pi/2} \Sigma_p \geqslant \Bbbk + 1$.

(a)⇒(b). Assume $(\xi^0, \xi^1, \dots, \xi^\Bbbk)$ is an array of directions in Σ_p, such that $\measuredangle(\xi^i, \xi^j) > \frac{\pi}{2}$ if $i \neq j$.

Without loss of generality, we may assume that each direction ξ^i is geodesic; that is, for each i there is a geodesic γ^i in $\mathcal{L}$ such that $\gamma^i(0) = p$ and $\xi^i = (\gamma^i)^+(0)$. From the definition of angle, it follows that for sufficiently small $\varepsilon > 0$ the array of points $a^i = \gamma^i(\varepsilon)$ satisfies (b). □

15.4. Corollary. *Let $\mathcal{L}$ be a complete length* CBB *space and $\Bbbk \in \mathbb{Z}_{\geqslant 0}$. Then the set of all points in $\mathcal{L}$ with rank $\geqslant \Bbbk$ is open.*

Proof. Given an array of points $\boldsymbol{a} = (a^0, \dots, a^\Bbbk)$ in $\mathcal{L}$, consider the set $\Omega_{\boldsymbol{a}}$ of all points $p \in \mathcal{L}$ such that array $\boldsymbol{a}$ is κ-strutting for a point p. Clearly $\Omega_{\boldsymbol{a}}$ is open.

According to Proposition 15.3, the set of points in $\mathcal{L}$ with rank $\geqslant \Bbbk$ can be presented as

$$\bigcup_{\boldsymbol{a}} \Omega_{\boldsymbol{a}},$$

where the union is taken over all $\Bbbk$-arrays $\boldsymbol{a}$ of points in $\mathcal{L}$. Hence the result. □

B Right-inverse theorem

15.5. Right-inverse theorem. *Suppose $\mathcal{L}$ is a complete length* CBB(κ) *space, $p, b \in \mathcal{L}$, and $\boldsymbol{a} = (a^1, \dots, a^\Bbbk)$ is a point array in $\mathcal{L}$.*

Assume that $(b, a^1, a^2, \dots, a^\Bbbk)$ is κ-strutting for p. Then the distance map $\operatorname{dist}_{\boldsymbol{a}} : \mathcal{L} \to \mathbb{R}^\Bbbk$ has a right inverse defined in a neighborhood of $\operatorname{dist}_{\boldsymbol{a}} p \in \mathbb{R}^\Bbbk$; that is, there is a submap $\Phi\colon \mathbb{R}^\Bbbk \hookrightarrow \mathcal{L}$ such that $\operatorname{Dom}\Phi \ni \operatorname{dist}_{\boldsymbol{a}} p$ and $\operatorname{dist}_{\boldsymbol{a}}[\Phi(\boldsymbol{x})] = \boldsymbol{x}$ for any $\boldsymbol{x} \in \operatorname{Dom}\Phi$. Moreover, *a) The map* Φ *can be chosen to be* $C^{\frac{1}{2}}$*-continuous (that is, Hölder continuous with exponent* $\frac{1}{2}$*) and such that*

$$\Phi(\operatorname{dist}_{\boldsymbol{a}} p) = p.$$

b) The distance map $\operatorname{dist}_{\boldsymbol{a}} : \mathcal{L} \to \mathbb{R}^{k}$ *is locally co-Lipschitz (in particular, open) in a neighborhood of* p.

Part b of the theorem is closely related to [34, Theorem 5.4] by Burago, Gromov and Perelman, but the proof presented here is different. Yet another proof can be built on [90, Proposition 4.3] by Lytchak.

Proof. Fix $\varepsilon, r, \lambda > 0$ such that the following conditions hold:

(i) Each distance function $\operatorname{dist}_{a^i}$ and dist_b is $\frac{\lambda}{2}$-concave in $\mathrm{B}(p,r)$.

(ii) For any $q \in \mathrm{B}(p,r)$, we have $\measuredangle^{\kappa}(q\,{}^{a^i}_{a^j}) > \frac{\pi}{2} + \varepsilon$ for all $i \neq j$ and $\measuredangle^{\kappa}(q\,{}^{b}_{a^i}) > \frac{\pi}{2} + \varepsilon$ for all i. In addition, $\varepsilon < \frac{1}{10}$.

Given $\boldsymbol{x} = (x^1, x^2, \dots, x^k) \in \mathbb{R}^k$, consider the function $f_{\boldsymbol{x}} \colon \mathcal{L} \to \mathbb{R}$ defined by

$$f_{\boldsymbol{x}} = \min_i\{h^i_{\boldsymbol{x}}\} + \varepsilon \cdot \operatorname{dist}_b,$$

where $h^i_{\boldsymbol{x}}(q) = \min\{0, |a^i - q| - x^i\}$. Note that for any $\boldsymbol{x} \in \mathbb{R}^k$, the function $f_{\boldsymbol{x}}$ is $(1+\varepsilon)$-Lipschitz and λ-concave in $\mathrm{B}(p,r)$. Denote by $\alpha_{\boldsymbol{x}}(t)$ the $f_{\boldsymbol{x}}$-gradient curve (see Chapter 16) that starts at p.

❶ *If for some* $\boldsymbol{x} \in \mathbb{R}^k$ *and* $t_0 \leqslant \frac{r}{2}$ *we have* $|\operatorname{dist}_{\boldsymbol{a}} p - \boldsymbol{x}| \leqslant \frac{\varepsilon^2}{10}\cdot t_0$*, then* $\operatorname{dist}_{\boldsymbol{a}}[\alpha_{\boldsymbol{x}}(t_0)] = \boldsymbol{x}$*.*

First note that Claim ❶ follows if for any $q \in \mathrm{B}(p,r)$, we have

(i) $(\boldsymbol{d}_q \operatorname{dist}_{a^i})(\nabla_q f_{\boldsymbol{x}}) < -\frac{1}{10}\cdot\varepsilon^2$ if $|a^i - q| > x^i$ and

(ii) $(\boldsymbol{d}_q \operatorname{dist}_{a^i})(\nabla_q f_{\boldsymbol{x}}) > \frac{1}{10}\cdot\varepsilon^2$ if

$$|a^i - q| - x^i = \min_j\{|a^j - q| - x^j\} < 0.$$

Indeed, since $t_0 \leqslant \frac{r}{2}$, then $\alpha_{\boldsymbol{x}}(t) \in \mathrm{B}(p,r)$ for all $t \in [0, t_0]$. Consider the following real-to-real functions:

❷
$$\begin{aligned}\varphi(t) &:= \max_i\{|a^i - \alpha_{\boldsymbol{x}}(t)| - x^i\},\\ \psi(t) &:= \min_i\{|a^i - \alpha_{\boldsymbol{x}}(t)| - x^i\}.\end{aligned}$$

Then from (i), we have $\varphi^+ < -\frac{1}{10}\cdot\varepsilon^2$ if $\varphi > 0$ and $t \in [0, t_0]$. Similarly, from (ii), we have $\psi^+ > \frac{1}{10}\cdot\varepsilon^2$ if $\psi < 0$ and $t \in [0, t_0]$. Since $|\operatorname{dist}_{\boldsymbol{a}} p - \boldsymbol{x}| \leqslant \leqslant \frac{\varepsilon^2}{10}\cdot t_0$, it follows that $\varphi(0) \leqslant \frac{\varepsilon^2}{10}\cdot t_0$ and $\psi(0) \geqslant -\frac{\varepsilon^2}{10}\cdot t_0$. Thus $\varphi(t_0) \leqslant 0$ and $\psi(t_0) \geqslant 0$. On the other hand, from ❷ we have $\varphi(t_0) \geqslant \psi(t_0)$. That is, $\varphi(t_0) = \psi(t_0) = 0$; hence Claim ❶ follows.

Thus, to prove Claim ❶, it remains to prove (i) and (ii). First let us prove it assuming that $\mathcal{L}$ is geodesic.

Note that

❸ $$(\boldsymbol{d}_q \operatorname{dist}_b)(\uparrow_{[qa^i]}) \leqslant \cos \tilde\measuredangle^\kappa(q\,{}^{b}_{a^j}) < -\tfrac{\varepsilon}{2}$$

for all i, and

❹ $$(\boldsymbol{d}_q \operatorname{dist}_{a^j})(\uparrow_{[qa^i]}) \leqslant \cos \tilde\measuredangle^\kappa(q\,{}^{a^i}_{a^j}) < -\tfrac{\varepsilon}{2}$$

for all $j \neq i$. Further, ❹ implies

❺ $$(\boldsymbol{d}_q h^j_{\boldsymbol{x}})(\uparrow_{[qa^i]}) \leqslant 0.$$

for all $i \neq j$. The assumption in (i) implies

$$\boldsymbol{d}_q f_{\boldsymbol{x}} = \min_{j\neq i}\{\boldsymbol{d}_q h^j_{\boldsymbol{x}}\} + \varepsilon\cdot(\boldsymbol{d}_q \operatorname{dist}_b).$$

Thus

$$\begin{aligned}
-(\boldsymbol{d}_q \operatorname{dist}_{a^i})(\nabla_q f_{\boldsymbol{x}}) &\geqslant \langle \uparrow_{[qa^i]}, \nabla_q f_{\boldsymbol{x}}\rangle \geqslant \\
&\geqslant (\boldsymbol{d}_q f_{\boldsymbol{x}})(\uparrow_{[qa^i]}) = \\
&= \min_{i\neq j}\{(\boldsymbol{d}_q h^i_{\boldsymbol{x}})(\uparrow_{[qa^i]})\} + \varepsilon\cdot(\boldsymbol{d}_q \operatorname{dist}_b)(\uparrow_{[qa^i]}).
\end{aligned}$$

Therefore (i) follows from ❸ and ❺.

The assumption in (ii) implies that $f_{\boldsymbol{x}}(q) = h^i_{\boldsymbol{x}}(q) + \varepsilon\cdot \operatorname{dist}_b$ and

$$\boldsymbol{d}_q f_{\boldsymbol{x}} \leqslant \boldsymbol{d}_q \operatorname{dist}_{a^i} + \varepsilon\cdot(\boldsymbol{d}_p \operatorname{dist}_b).$$

Therefore,

$$\begin{aligned}
(\boldsymbol{d}_q \operatorname{dist}_{a^i})(\nabla_q f_{\boldsymbol{x}}) &\geqslant \boldsymbol{d}_q f_{\boldsymbol{x}}(\nabla_q f_{\boldsymbol{x}}) \geqslant \\
&\geqslant \Big[(\boldsymbol{d}_q f_{\boldsymbol{x}})(\uparrow_{[qb]})\Big]^2 \geqslant \\
&\geqslant \Big[\min_i\{\cos \tilde\measuredangle^\kappa(q\,{}^{b}_{a^i})\} - \varepsilon^2\Big]^2.
\end{aligned}$$

Thus (ii) follows from ❸, since $\varepsilon < \frac{1}{10}$.

Therefore ❶ holds if $\mathcal{L}$ is geodesic. If $\mathcal{L}$ is not geodesic, perform the above estimate in $\mathcal{L}^\omega$, the ultrapower of $\mathcal{L}$. (Recall that according to 3.6, $\mathcal{L}^\omega$ is geodesic.) This completes the proof of ❶. △

Set $t_0(\boldsymbol{x}) = \frac{10}{\varepsilon^2}\cdot|\operatorname{dist}_{\boldsymbol{a}} p - \boldsymbol{x}|$, giving equality in ❶. Define the submap Φ by

$$\Phi\colon \boldsymbol{x} \mapsto \alpha_{\boldsymbol{x}} \circ t_0(\boldsymbol{x}), \quad \operatorname{Dom}\Phi = \mathrm{B}(\operatorname{dist}_{\boldsymbol{a}} p, \tfrac{\varepsilon^2\cdot r}{20}) \subset \mathbb{R}^{\Bbbk}.$$

It follows from Claim ❶ that $\operatorname{dist}_{\boldsymbol{a}}[\Phi(\boldsymbol{x})] = \boldsymbol{x}$ for any $\boldsymbol{x} \in \operatorname{Dom}\Phi$.

Clearly $t_0(p) = 0$; thus $\Phi(\operatorname{dist}_{\boldsymbol{a}} p) = p$. Further, by construction of $f_{\boldsymbol{x}}$,

$$|f_{\boldsymbol{x}}(q) - f_{\boldsymbol{y}}(q)| \leqslant |\boldsymbol{x} - \boldsymbol{y}|,$$

for any $q \in \mathcal{L}$. Therefore, according to Lemma 16.13, Φ is $C^{\frac{1}{2}}$-continuous. Thus (a).

Further, note that

❻ $$|p - \Phi(\boldsymbol{x})| \leqslant (1 + \varepsilon)\cdot t_0(\boldsymbol{x}) \leqslant \tfrac{11}{\varepsilon^2}\cdot|\operatorname{dist}_{\boldsymbol{a}} p - \boldsymbol{x}|$$

holds.

The above construction may be repeated for any $p' \in \mathrm{B}(p, \frac{r}{4})$, $\varepsilon' = \varepsilon$, and $r' = \frac{r}{2}$. The inequality ❻ for the resulting map Φ' implies that for any $p', q \in \mathrm{B}(p, \frac{r}{4})$ there is $q' \in \mathcal{L}$ such that $\Phi'(q) = \Phi'(q')$ and

$$|p' - q'| \leqslant \tfrac{11}{\varepsilon^2}\cdot|\operatorname{dist}_{\boldsymbol{a}} p' - \boldsymbol{x}|.$$

That is, the distance map $\operatorname{dist}_{\boldsymbol{a}}$ is locally $\frac{11}{\varepsilon^2}$-co-Lipschitz in $\mathrm{B}(p, \frac{r}{4})$. □

C Dimension theorem for CBB

The following theorem is the main result of this section.

15.6. Theorem. *Let $\mathcal{L}$ be a complete length $\mathrm{CBB}(\kappa)$ space, $q \in \mathcal{L}$, $R > 0$ and $m \in \mathbb{Z}_{\geqslant 0}$. Then the following statements are equivalent:*

A) $\operatorname{LinDim}\mathcal{L} \geqslant m$.

B) There is a point $p \in \mathcal{L}$ that admits a κ-strutting array $(b, a^1, \dots \dots, a^m) \in \mathcal{L}^{m+1}$.

C) Let Euk^m be the set of all points $p \in \mathcal{L}$ such that there is a distance-preserving embedding $\mathbb{E}^m \hookrightarrow \mathrm{T}_p$ that preserves the cone structure (see Section 6E). Then Euk^m contains a dense G-delta set in $\mathcal{L}$.

D) There is a $C^{\frac{1}{2}}$-embedding; that is, a bi-Hölder embedding with exponent $\frac{1}{2}$,

$$\overline{\mathrm{B}}[1]_{\mathbb{E}^m} \hookrightarrow \mathrm{B}(q, R).$$

E)

$$\operatorname{pack}_\varepsilon \mathrm{B}(q, R) > \frac{c}{\varepsilon^m}$$

for some fixed $c > 0$ and any $\varepsilon > 0$.

In particular:

(i) If $\operatorname{LinDim}\mathcal{L} = \infty$, then all the statements (C), (D), and (E) are satisfied for all $m \in \mathbb{Z}_{\geqslant 0}$.

(ii) If the statement (D) or (E) is satisfied for some choice of $q \in \in \mathcal{L}$ *and* $R > 0$*, then it also is satisfied for any other choice of* q *and* R.

For finite-dimensional spaces, Theorem 15.13 gives a stronger version of the theorem above.

The proof of the above theorem with the exception of statement (D) was given in [111]. At that time, it was not known whether for any complete length CBB(κ) space $\mathcal{L}$,

$$\operatorname{LinDim}\mathcal{L} = \infty \quad \Rightarrow \quad \operatorname{TopDim}\mathcal{L} = \infty.$$

The latter implication was proved by Grigory Perelman and the third author [101]; it was done by combining an idea of Conrad Plaut with the technique of gradient flow. The statement 15.6D is somewhat stronger.

To prove Theorem 15.6 we will need the following three propositions.

15.7. Proposition. *Let* $\mathcal{L}$ *be a complete length* CBB(κ) *space and* $p \in \in \mathcal{L}$. *Assume there is a distance-preserving embedding* $\iota\colon \mathbb{E}^m \hookrightarrow \mathrm{T}_p\mathcal{L}$ *that preserves the cone structure. Then either*

a) $\operatorname{Im}\iota = \mathrm{T}_p\mathcal{L}$, *or*

b) there is a point p' *arbitrarily close to* p *such that there is a distance-preserving embedding* $\iota'\colon \mathbb{E}^m \hookrightarrow \mathrm{T}_{p'}\mathcal{L}$ *that preserves the cone structure.*

Proof. Assume $\iota(\mathbb{E}^m)$ is a proper subset of $\mathrm{T}_p\mathcal{L}$. Equivalently, there is a direction $\xi \in \Sigma_p\backslash\iota(\mathbb{S}^{m-1})$, where $\mathbb{S}^{m-1} \subset \mathbb{E}^m$ is the unit sphere.

Fix $\varepsilon > 0$ so that $\measuredangle(\xi, \sigma) > \varepsilon$ for any $\sigma \in \iota(\mathbb{S}^{m-1})$. Choose a maximal ε-packing in $\iota(\mathbb{S}^{m-1})$; that is, an array $(\zeta^1, \zeta^2, \dots, \zeta^n)$ of directions in $\iota(\mathbb{S}^{m-1})$ such that $n = \operatorname{pack}_\varepsilon \mathbb{S}^{m-1}$ and $\measuredangle(\zeta^i, \zeta^j) > \varepsilon$ for any $i \ne j$.

Choose an array $(x, z^1, z^2, \dots, z^n)$ of points in $\mathcal{L}$ such that $\uparrow_{[px]} \approx \xi$, $\uparrow_{[pz^i]} \approx \zeta^i$; here we write "≈" for "sufficiently close". We can choose this array so $\tilde\measuredangle^\kappa(p\,{}^x_{z^i}) > \varepsilon$ for all i and $\tilde\measuredangle^\kappa(p\,{}^{z^i}_{z^j}) > \varepsilon$ for all $i \ne j$. Applying Corollary 13.37, there is a point p' arbitrarily close to p such that all directions $\uparrow_{[p'x]}, \uparrow_{[p'z^1]}, \uparrow_{[p'z^2]}, \dots, \uparrow_{[p'z^n]}$ belong to an isometric copy of $\mathbb{S}^{\hbar-1}$ in $\Sigma_{p'}$. In addition, we may assume that $\tilde\measuredangle^\kappa(p'\,{}^x_{z^i}) > \varepsilon$ and $\tilde\measuredangle^\kappa(p'\,{}^{z^i}_{z^j}) > \varepsilon$. From the hinge comparison (8.14c), $\measuredangle(\uparrow_{[p'x]}, \uparrow_{[p'z^i]}) > \varepsilon$ and $\measuredangle(\uparrow_{[p'z^i]}, \uparrow_{[p'z^j]}) > \varepsilon$; that is,

$$\operatorname{pack}_\varepsilon \mathbb{S}^{\hbar-1} \geqslant n + 1 > \operatorname{pack}_\varepsilon \mathbb{S}^{m-1}.$$

Hence $\hbar > m$. □

15.8. Proposition. *Let $\mathcal{L}$ be a complete length* CBB(κ) *space. Then for any two points $p, \bar{p} \in \mathcal{L}$ and any $R, \bar{R} > 0$, there is a constant $\delta = \delta(\kappa, R, \bar{R}, |p - \bar{p}|) > 0$ such that*

$$\operatorname{pack}_{\delta\cdot\varepsilon} \mathrm{B}(\bar{p}, \bar{R}) \geqslant \operatorname{pack}_{\varepsilon} \mathrm{B}(p, R).$$

Proof. According to 8.32, we can assume that $\kappa \leqslant 0$.

Let $n = \operatorname{pack}_\varepsilon \mathrm{B}(p, R)$ and $\{x^1, \dots, x^n\}$ be a maximal ε-packing in $\mathrm{B}(p, R)$; that is, $|x^i - x^j| > \varepsilon$ for all $i \neq j$. Without loss of generality, we may assume the x^i are in $Str(\bar{p})$. Thus, for each i there is a unique geodesic $[\bar{p}x^i]$ (see 8.11). Choose a factor $1 > s > 0$ so that $\bar{R} > s{\cdot}(|p - \bar{p}| + R)$. For each i, take $\bar{x}^i \in [\bar{p}x^i]$ so that $|\bar{p} - \bar{x}^i| = s{\cdot}(|p - x^i|)$. From 8.17$a$,

$$\tilde{\measuredangle}^\kappa\left(\bar{p}\,{}^{\bar{x}^i}_{\bar{x}^j}\right) \geqslant \tilde{\measuredangle}^\kappa(\bar{p}\,{}^{x^i}_{x^j}).$$

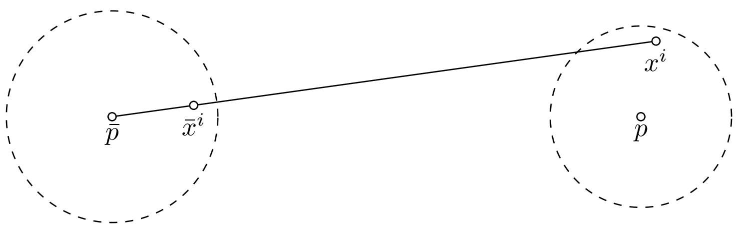

Applying the cosine law gives a constant $\delta = \delta(\kappa, R, \bar{R}, |p - \bar{p}|) > 0$ such that

$$|\bar{x}^i - \bar{x}^j| > \delta{\cdot}(|x^i - x^j|) > \delta{\cdot}\varepsilon$$

for all $i \neq j$. Hence the statement follows. □

15.9. Proposition. *Let $\mathcal{L}$ be a complete length* CBB(κ) *space, $r < \varpi^\kappa$ and $p \in \mathcal{L}$. Assume that*

❶ $$\operatorname{pack}_\varepsilon \mathrm{B}(p, r) > \operatorname{pack}_\varepsilon \overline{\mathrm{B}}[r]_{\mathbb{M}^m(\kappa)}$$

for some $\varepsilon > 0$. Then there is a G-delta set $A \subset \mathcal{L}$ that is dense in a neighborhood of p and such that $\dim \operatorname{Lin}_q > m$ for any $q \in A$.

Proof. Choose a maximal ε-packing in $\mathrm{B}(p, r)$, that is, an array $(x^1, x^2, \dots, x^n)$ of points in $\mathrm{B}(p, r)$ such that $n = \operatorname{pack}_\varepsilon \mathrm{B}(p, r)$ and $|x^i - x^j| > \varepsilon$ for any $i \neq j$. Choose a neighborhood $\Omega \ni p$ such that $|q - x^i| < r$ for any $q \in \Omega$ and all i. Let

$$A = \Omega \cap \operatorname{Str}(x^1, x^2, \dots, x^n).$$

According to Theorem 8.11, A is a G-delta set that is dense in Ω.

Assume $\hbar = \dim \mathrm{Lin}_q \leqslant m$ for some $q \in A$. Consider an array $(v^1, v^2, \dots, v^n)$ of vectors in Lin_q, where $v^i = \log[qx^i]$. Clearly

$$|v^i| = |q - x^i| < r,$$

and from the hinge comparison (8.14c) we have

$$\tilde{\curlyvee}^{\kappa}[0\,{}^{v^i}_{v^j}] \geqslant |x^i - x^j| > \varepsilon.$$

Note that the ball $\mathrm{B}(0, r)_{\mathrm{Lin}_q}$ equipped with the metric $\rho(v, w) = \tilde{\curlyvee}^{\kappa}[0\,{}^{v}_{w}]$ is isometric to $\overline{\mathrm{B}}[r]_{\mathbb{M}^{\hbar}(\kappa)}$. Thus

$$\mathrm{pack}_{\varepsilon}\,\overline{\mathrm{B}}[r]_{\mathbb{M}^{\hbar}(\kappa)} \geqslant \mathrm{pack}_{\varepsilon}\,\mathrm{B}(p, r),$$

which contradicts $\hbar \leqslant m$ and ❶. □

The proof of Theorem 15.6 is essentially done in 15.7, 15.8, 15.9, 15.10, 15.5; now we assemble the proof from these parts.

We will prove the implications

$$(C) \Rightarrow (A) \Rightarrow (B) \Rightarrow (E) \Rightarrow (C) \Rightarrow (D) \Rightarrow (E).$$

Proof of 15.6. The implication $(C)\Rightarrow(A)$ is trivial.

(A)⇒(B). Choose a point $p \in \mathcal{L}$ such that $\dim \mathrm{Lin}_p \geqslant m$. Clearly one can choose an array $(\xi^0, \xi^1, \dots, \xi^m)$ of directions in Lin_p such that $\measuredangle(\xi^i, \xi^j) > \frac{\pi}{2}$ for all $i \neq j$. Choose an array $(x^0, x^1, \dots, x^m)$ of points in $\mathcal{L}$ such that each $\uparrow_{[px^i]}$ is sufficiently close to ξ^i; in particular, we have $\measuredangle[p\,{}^{x^i}_{x^j}] > \frac{\pi}{2}$. Choose points $a^i \in\,]px^i]$ sufficiently close to p. This can be done so that each $\tilde{\measuredangle}^{\kappa}(p\,{}^{a^i}_{a^j})$ is arbitrarily close to $\measuredangle[p\,{}^{a^i}_{a^j}]$, in particular $\tilde{\measuredangle}^{\kappa}(p\,{}^{a^i}_{a^j}) > \frac{\pi}{2}$. Finally, set $b = a^0$.

(B)⇒(E). Let $p \in \mathcal{L}$ be a point that admits a κ-strutting array $(b, a^1, \dots$ $\dots, a^m)$ of points in $\mathcal{L}$. The right-inverse mapping map theorem (15.5b) implies that the distance map $\mathrm{dist}_{\boldsymbol{a}} : \mathcal{L} \to \mathbb{R}^m$,

$$\mathrm{dist}_{\boldsymbol{a}} : x \mapsto (|a^1 - x|, |a^2 - x|, \dots, |a^n - x|),$$

is open in a neighborhood of p. Since the distance map $\mathrm{dist}_{\boldsymbol{a}}$ is Lipschitz, for any $r > 0$, there is $c > 0$ such that

$$\mathrm{pack}_{\varepsilon}\,\mathrm{B}(p, r) > \frac{c}{\varepsilon^m}.$$

Applying 15.8, we get a similar inequality for any other ball in $\mathcal{L}$; that is, for any $q \in \mathcal{L}$ and $R > 0$, there is $c' > 0$ such that

$$\mathrm{pack}_{\varepsilon}\,\mathrm{B}(q, R) > \frac{c'}{\varepsilon^m}.$$

(E)⇒(C). Note that for any $q' \in \mathcal{L}$ and $R' > |q - q'| + R$ we have

$$\begin{aligned}\operatorname{pack}_\varepsilon \mathrm{B}(q', R') &\geqslant \operatorname{pack}_\varepsilon \mathrm{B}(q, R) \geqslant \\ &\geqslant \frac{c}{\varepsilon^m} > \\ &> \operatorname{pack}_\varepsilon \overline{\mathrm{B}}[R']_{\mathbb{M}^{m-1}(\kappa)}\end{aligned}$$

for all sufficiently small $\varepsilon > 0$. Applying 15.9, Euk^m contains a G-delta set that is dense in a neighborhood of any point $q' \in \mathcal{L}$.

(C)⇒(D). Since Euk^m contains a dense G-delta set in $\mathcal{L}$, we can choose $p \in \mathrm{B}(q, R)$ with a distance-preserving cone embedding $\iota\colon \mathbb{E}^m \hookrightarrow \mathrm{T}_p$.

Repeating the construction in (A)⇒(B), we get a κ-strutting array $(p, a^1, \dots, a^m)$ for p.

Applying the right-inverse theorem (15.5), we obtain a $C^{\frac{1}{2}}$-submap

$$\Phi\colon \mathbb{R}^m \mathrel{\circ\!\!\to} \mathrm{B}(q, R)$$

that is a right inverse for $\operatorname{dist}_{\boldsymbol{a}} : \mathcal{L} \to \mathbb{R}^m$ and such that $\Phi(\operatorname{dist}_{\boldsymbol{a}} p) = p$. In particular, Φ is a $C^{\frac{1}{2}}$-embedding of $\operatorname{Dom}\Phi$.

(D)⇒(E). This proof is valid for general metric spaces; it is based on general relations between topological dimension, Hausdorff measure and $\operatorname{pack}_\varepsilon$.

Let $W \subset \mathrm{B}(q, R)$ be the image of the embedding Φ. Since $\operatorname{TopDim} W = m$, Szpilrajn's theorem (7.5) implies that

$$\operatorname{HausMes}_m W > 0.$$

Given $\varepsilon > 0$, consider a maximal ε-packing of W, that is, an array $(x^1, x^2, \dots, x^n)$ of points in W such that $n = \operatorname{pack}_\varepsilon W$ and $|x^i - x^j| > \varepsilon$ for all $i \neq j$. Note that W is covered by balls $\mathrm{B}(x^i, 2\cdot\varepsilon)$.

By the definition of Hausdorff measure,

$$\operatorname{pack}_\varepsilon W \geqslant \frac{c}{\varepsilon^m} \cdot \operatorname{HausMes}_m W$$

for a fixed constant $c > 0$ and all small $\varepsilon > 0$. Hence (E) follows. □

D Inverse function theorem

15.10. Inverse function theorem. *Let $\mathcal{L}$ be an m-dimensional complete length* $\mathrm{CBB}(\kappa)$ *space and $p, b, a^1, a^2, \dots, a^m \in \mathcal{L}$.*

Assume that the point array $\boldsymbol{a} = (b, a^1, \dots, a^m)$ is κ-strutting for p. Then there are $R > 0$ and $\varepsilon > 0$ such that:

a) For all $i\ne j$ and any $q\in\mathrm{B}(p,R)$ we have

$$\tilde\measuredangle^\kappa(q\,{}^{a^i}_{a^j})>\tfrac{\pi}{2}+\varepsilon\quad\textit{and}\quad\tilde\measuredangle^\kappa(q\,{}^{b}_{a^i})>\tfrac{\pi}{2}.$$

b) The restriction of the distance map

$$\mathrm{dist}_{\boldsymbol{a}}\colon x\mapsto(|a^1-x|,\dots,|a^m-x|)$$

to the ball $\mathrm{B}(p,R)$ is an open $[\varepsilon,\sqrt{m}]$-bi-Lipschitz embedding $\mathrm{B}(p,R)\hookrightarrow\mathbb{R}^m$.

c) The value R depends only on

$$\kappa,\quad|p-a^i|,\quad|a^i-a^j|,\quad|b-a^i|$$

for all i and j.

15.11. Definition. *Suppose $\mathcal{L}$ is an m-dimensional complete length $\mathrm{CBB}(\kappa)$ space. If a point array $(b,a^1,a^2,\dots,a^m)$ and the value R satisfy the conditions in Theorem 15.10, then the restriction $\boldsymbol{x}=\mathrm{dist}_{\boldsymbol{a}}|_{\mathrm{B}(p,R)}$ is called a distance chart, the restrictions $x^i=\mathrm{dist}_{a^i}|_{\mathrm{B}(p,R)}$ are called coordinates, and the restriction $y=\mathrm{dist}_b|_{\mathrm{B}(p,R)}$ is called the strut of the distance chart.*

The proof of Theorem 15.10 will require the following lemma.

15.12. Lemma. *Suppose $\mathcal{L}$ is an m-dimensional complete length $\mathrm{CBB}(\kappa)$ space and $p\in\mathcal{L}$. Assume for the directions $\xi,\zeta^1,\zeta^2,\dots,\zeta^{\hbar}\in\Sigma_p$ the following conditions hold:*

a) $\measuredangle(\xi,\zeta^i)>\frac{\pi}{2}-\varepsilon$ for all i,

b) $\measuredangle(\zeta^i,\zeta^j)>\frac{\pi}{2}+\varepsilon$ for all $i\ne j$.

Then $\hbar\leqslant m$.

Proof. Without loss of generality, we can assume that all $\xi,\zeta^1,\zeta^2,\dots,\zeta^{\hbar}$ are geodesic directions; let $\xi=\uparrow_{[px]}$ and $\zeta^i=\uparrow_{[pz^i]}$ for all i. Fix a small $r>0$, and let $\bar x\in\,]px]$ and $\bar z^i\in\,]pz^i]$ be points such that

$$|p-\bar x|=|p-\bar z^1|=\dots=|p-\bar z^{\hbar}|=r.$$

From the definition of angle, if r is sufficiently small we have

$\diamond$ $\tilde\measuredangle^\kappa(p\,{}^{\bar x}_{\bar z^i})>\frac{\pi}{2}-\varepsilon$ for all i, and $\tilde\measuredangle^\kappa(p\,{}^{\bar z^i}_{\bar z^j})>\frac{\pi}{2}+\varepsilon$ for all $i\ne j$.

Choose a point $p'\in\mathrm{Str}(\bar x,\bar z^1,\bar z^2,\dots,\bar z^{\hbar})$ sufficiently close to p that the above conditions still hold for p'; that is,

❶ $\tilde\measuredangle^\kappa(p'\,{}^{\bar x}_{\bar z^i})>\frac{\pi}{2}-\varepsilon$ *for all i, and* $\tilde\measuredangle^\kappa(p'\,{}^{\bar z^i}_{\bar z^j})>\frac{\pi}{2}+\varepsilon$ *for all $i\ne j$.*

Set $\acute\xi=\uparrow_{[p'\bar x]}$ and $\acute\zeta^i=\uparrow_{[p'\bar z^i]}$ for each i. By the hinge comparison (8.14c), ❷ $\measuredangle(\acute{\xi},\acute{\zeta}^i) > \frac{\pi}{2} - \varepsilon$ *for all* i, *and* $\measuredangle(\acute{\zeta}^i,\acute{\zeta}^j) > \frac{\pi}{2} + \varepsilon$ *for all* $i \neq j$.

According to Corollary 13.37, all directions $\acute{\xi},\acute{\zeta}^1,\acute{\zeta}^2,\dots,\acute{\zeta}^{\hbar}$ lie in an isometric copy of the standard n-sphere in $\Sigma_{p'}$. Clearly $n \leqslant m-1$. Thus it remains to prove the following claim, which is a partial case of the lemma.

❸ *If* $\xi,\zeta^1,\zeta^2,\dots,\zeta^{\hbar} \in \mathbb{S}^{m-1}$, $|\xi-\zeta^i| > \frac{\pi}{2}-\varepsilon$ *for all* i, *and* $|\zeta^i-\zeta^j| > \frac{\pi}{2}+\varepsilon$ *for all* $i \neq j$, *then* $\hbar \leqslant m$.

For each i, let $\bar{\zeta}^i$ be the closest point to ζ^i in $\Xi = \mathbb{S}^{m-1}\backslash \mathrm{B}(\xi,\frac{\pi}{2}) \stackrel{iso}{=\!=} \mathbb{S}^{m-1}_+$ (if $\zeta \in \Xi$, then $\bar{\zeta}^i = \zeta^i$). By straightforward calculations, we have

$$|\bar{\zeta}^i - \bar{\zeta}^j| \geqslant |\zeta^i - \zeta^j| - \varepsilon > \tfrac{\pi}{2}.$$

Thus it is sufficient to show the following claim:

❹ $\operatorname{pack}_{\frac{\pi}{2}} \mathbb{S}^{m-1}_+ = m$.

Clearly, $\operatorname{pack}_{\frac{\pi}{2}} \mathbb{S}^{m-1}_+ \geqslant m$.

The opposite inequality is proved by induction on m. The base case $m=1$ is obvious. Assume $(\bar{\zeta}^1,\bar{\zeta}^2,\dots,\bar{\zeta}^{\hbar})$ is an array of points in $\mathbb{S}^{m-1}_+$ with $|\bar{\zeta}^i - \bar{\zeta}^j| > \frac{\pi}{2}$. Without loss of generality, we can also assume that $\bar{\zeta}^{\hbar} \in \partial\mathbb{S}^{m-1}_+$. For each $i < \hbar$, let $\check{\zeta}^i = \uparrow_{[\bar{\zeta}^{\hbar}\bar{\zeta}^i]} \in \Sigma_{\bar{\zeta}^{\hbar}}\mathbb{S}^{m-1}_+ \stackrel{iso}{=\!=} \mathbb{S}^{m-2}_+$. By the hinge comparison (8.14*c*), $\measuredangle(\check{\zeta}^i,\check{\zeta}^j) > \frac{\pi}{2}$ for all $i < j < \hbar$. Thus from the induction hypothesis we have $\hbar - 1 \leqslant m-1$. □

Proof of 15.10; (a). Fix $\varepsilon > 0$ such that $\tilde{\measuredangle}^{\kappa}(p\,{}^{a^i}_{a^j}) > \frac{\pi}{2}+\varepsilon$ and $\tilde{\measuredangle}^{\kappa}(p\,{}^{b}_{a^i}) > \frac{\pi}{2}+\varepsilon$ for all $i \neq j$. Choose $R > 0$ sufficiently small that $\tilde{\measuredangle}^{\kappa}(q\,{}^{a^i}_{a^j}) > \frac{\pi}{2}+\varepsilon$ and $\tilde{\measuredangle}^{\kappa}(q\,{}^{b}_{a^i}) > \frac{\pi}{2}+\varepsilon$ for all $i \neq j$ and any $q \in \mathrm{B}(p,R)$. Clearly, (*a*) holds for $\mathrm{B}(p,R)$.

(b). Note that the distance map $\operatorname{dist}_{\boldsymbol{a}}$ is Lipschitz and its restriction $\operatorname{dist}_{\boldsymbol{a}}|_{\mathrm{B}(p,R)}$ is open; the latter follows from the right-inverse theorem (15.5*b*). Thus to prove (*b*), it is sufficient to show that

❺ $$\max_i\{\,||a^i-x|-|a^i-y||\,\} > \tfrac{\varepsilon}{2}\cdot|x-y|$$

for any $x,y \in \mathrm{B}(p,R)$.

According to Lemma 15.12,

$$\measuredangle[x\,{}^{y}_{b}] \leqslant \tfrac{\pi}{2}-\varepsilon \quad\text{or}\quad \measuredangle[x\,{}^{y}_{a^i}] \leqslant \tfrac{\pi}{2}-\varepsilon \quad\text{for some}\quad i.$$

In the latter case, since $|x-y| < 2\cdot R$ and R is small, the hinge comparison (8.14*c*) implies

❻ $$|a^i-x|-|a^i-y| > \tfrac{\varepsilon}{2}\cdot|x-y| \quad\text{for some}\quad i.$$

If $\measuredangle[x\,{}^{y}_{b}] \leqslant \frac{\pi}{2} - \varepsilon$, then switching x and y, we get

❼ $$|a^j - y| - |a^j - x| > \tfrac{\varepsilon}{2}\cdot|x - y| \quad \text{for some} \quad j.$$

Then ❻ and ❼ imply ❺.

Finally, part (c) follows since the angle $\tilde\measuredangle^\kappa(q\,{}^{a^i}_{a^j})$ depends continuously on κ, $|q - a^i|$, $|q - a^j|$ and $|a^i - a^j|$. □

E Finite dimensional CBB spaces

In this section we show that all reasonable notions of dimension coincide on the class of Alexandrov spaces with curvature bounded below.

First we prove a stronger version of Theorem 15.6 for the finite-dimensional case.

15.13. Theorem. *Suppose $\mathcal{L}$ is a complete length* CBB(κ) *space, m is a nonnegative integer, $0 < R \leqslant \varpi^\kappa$, and $q \in \mathcal{L}$. Then the following statements are equivalent:*

a) LinDim $\mathcal{L} = m$.

b) m is the maximal integer such that there is a point $p \in \mathcal{L}$ that admits a κ-strutting array $(b, a^1, \dots, a^m)$.

c) $\mathrm{T}_p \overset{iso}{=\!=} \mathbb{E}^m$ *for any point p in a dense G-delta set of $\mathcal{L}$.*

d) There is an open bi-Lipschitz embedding

$$\overline{\mathrm{B}}[1]_{\mathbb{E}^m} \hookrightarrow \mathrm{B}(q, R) \subset \mathcal{L}.$$

e) For any $\varepsilon > 0$,

$$\mathrm{pack}_\varepsilon \overline{\mathrm{B}}[R]_{\mathbb{M}^m(\kappa)} \geqslant \mathrm{pack}_\varepsilon \mathrm{B}(q, R).$$

moreover, there is $c = c(q, R) > 0$ such that

$$\mathrm{pack}_\varepsilon \mathrm{B}(q, R) > \frac{c}{\varepsilon^m}.$$

The above theorem was essentially proved in [34].

Using theorems 15.6 and 15.13, one can show that linear dimension is equal to many different types of dimension, such as small and big inductive dimension and upper and lower box-counting dimension (also known as Minkowski dimension), homological dimension and so on.

The next two corollaries follow from 15.13e.

15.14. Corollary. *Any m-dimensional complete length* CBB *space is proper and geodesic.*

15.15. Corollary. *Let $(\mathcal{L}_n)$ be a sequence of length* CBB(κ) *spaces and $\mathcal{L}_n \to \mathcal{L}_\omega$ as $n \to \omega$. Assume* $\operatorname{LinDim} L_n \leqslant m$ *for all n. Then* $\operatorname{LinDim} L_\omega \leqslant m$.

15.16. Corollary. *Let $\mathcal{L}$ be a complete length* CBB(κ) *space. Then for any open $\Omega \subset \mathcal{L}$, we have*

$$\operatorname{LinDim} \mathcal{L} = \operatorname{LinDim} \Omega = \operatorname{TopDim} \Omega = \operatorname{HausDim} \Omega,$$

where TopDim *and* HausDim *denote topological dimension (7.2) and Hausdorff dimension (7.1) respectively.*

In particular, $\mathcal{L}$ is dimension-homogeneous; that is, all open sets have the same linear dimension.

Proof of 15.16. The equality

$$\operatorname{LinDim} \mathcal{L} = \operatorname{LinDim} \Omega$$

follows from 15.6A&C.

If $\operatorname{LinDim} \mathcal{L} = \infty$, then applying 15.6$D$ for $\mathrm{B}(q, R) \subset \Omega$, we find that there is a compact subset $K \subset \Omega$ having an arbitrarily large $\operatorname{TopDim} K$. Therefore

$$\operatorname{TopDim} \Omega = \infty.$$

By Szpilrajn's theorem (7.5), $\operatorname{HausDim} K \geqslant \operatorname{TopDim} K$. Thus we also have

$$\operatorname{HausDim} \Omega = \infty.$$

If $\operatorname{LinDim} \mathcal{L} = m < \infty$, then the first inequality in 15.13e implies that

$$\operatorname{HausDim} \mathrm{B}(q, R) \leqslant m.$$

According to Corollary 15.14, $\mathcal{L}$ is proper and in particular has countable base. Thus applying Szpilrajn's theorem again, we have

$$\operatorname{TopDim} \Omega \leqslant \operatorname{HausDim} \Omega \leqslant m.$$

Finally, 15.13d implies that $m \leqslant \operatorname{TopDim} \Omega$. □

Proof of 15.13. The equivalence $(a)\Leftrightarrow(b)$ follows from 15.6.

(a)⇒(c). If $\operatorname{LinDim} \mathcal{L} = m$, then by Theorem 15.6, Euk^m contains a dense G-delta set in $\mathcal{L}$. From 15.7, it follows that T_p is isometric to $\mathbb{E}^m$ for any $p \in \mathrm{Euk}^m$.

(c)⇒(d). This is proved in exactly the same way as implication *(C)⇒(D)* of theorem 15.6, but applying the existence of a distance chart (15.10) instead of the right-inverse theorem (15.5).

(d)⇒(e). From (d), it follows that there is a point $p\in \mathrm{B}(q,R)$ and $r>0$ such that $\mathrm{B}(p,r)\subset\mathcal{L}$ is bi-Lipschitz homeomorphic to a bounded open set of $\mathbb{E}^m$. Thus there is $c>0$ such that

❶ $$\operatorname{pack}_\varepsilon \mathrm{B}(p,r) > \frac{c}{\varepsilon^m}.$$

Applying 15.8 shows that inequality ❶, with different constants, holds for any other ball, in particular for $\mathrm{B}(q,R)$.

Applying 15.9 gives the first inequality in (e).

(e)⇒(a). From Theorem 15.6, we have $\operatorname{LinDim}\mathcal{L}\geqslant m$. Applying Theorem 15.6 again, if $\operatorname{LinDim}\mathcal{L}\geqslant m+1$ then for some $c>0$ and any $\varepsilon>0$,

$$\operatorname{pack}_\varepsilon \mathrm{B}(q,R)\geqslant\frac{c}{\varepsilon^{m+1}}.$$

But

$$\frac{c'}{\varepsilon^m}\geqslant\operatorname{pack}_\varepsilon \mathrm{B}(q,R)$$

for any $\varepsilon>0$, a contradiction. □

The following exercise was suggested by Alexander Lytchak.

15.17. Exercise. *Suppose $\mathcal{L}$ is a complete length* CBB *space and $\Sigma_p\mathcal{L}$ is compact for any $p\in\mathcal{L}$. Prove that $\mathcal{L}$ is finite-dimensional.*

F One-dimensional CBB spaces

15.18. Theorem. *Let $\mathcal{L}$ be a one-dimensional complete length* CBB(κ) *space. Then $\mathcal{L}$ is isometric to a connected complete Riemannian* 1*-dimensional manifold with possibly non-empty boundary.*

Proof. Clearly $\mathcal{L}$ is connected. It remains to show the following:

❶ *For any point $p\in\mathcal{L}$ there is $\varepsilon>0$ such that* $\mathrm{B}(p,\varepsilon)$ *is isometric to either* $[0,\varepsilon)$ *or* $(-\varepsilon,\varepsilon)$.

First let us show:

❷ *If $p\in\,]xy[$ for some x, $y\in\mathcal{L}$ and $\varepsilon<\min\{|p-x|,|p-y|\}$, then* $\mathrm{B}(p,\varepsilon)\subset\,]xy[$. *In particular,* $\mathrm{B}(p,\varepsilon)\overset{iso}{=\!=}(-\varepsilon,\varepsilon)$.

Assume the contrary; that is, there is

$$z\in \mathrm{B}(p,\varepsilon)\backslash]xy[.$$

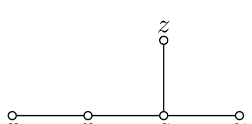

Consider a geodesic $[pz]$, and let $q\in[pz]\cap[xy]$ be the point that maximizes the distance $|p-q|$. At

q, we have three distinct directions: to x, y, and z. Moreover, $\measuredangle[q\,^x_y] = \pi$. Thus, according to Proposition 15.7, $\operatorname{LinDim}\mathcal{L} > 1$, a contradiction. $\triangle$

Now assume no geodesic includes p as a non endpoint. Since $\operatorname{LinDim}\mathcal{L} = 1$ there is a point $y \neq p$.

Fix a positive value $\varepsilon < |p - y|$. Let us show:

❸ $\mathrm{B}(p, \varepsilon) \subset [py]$; *in particular,* $\mathrm{B}(p, \varepsilon) \overset{iso}{=\!=} [0, \varepsilon)$.

Assume the contrary; let $z \in \mathrm{B}(p, \varepsilon)\backslash[py]$. Choose a point $w \in \,]py[$ such that

$$|p - w| + |p - z| < \varepsilon.$$

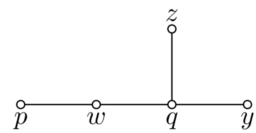

Consider geodesic $[wz]$, and let $q \in [py] \cap [wz]$ be the point that maximizes the distance $|w - q|$. Since no geodesic includes p as a non endpoint, we have $p \neq q$. As above, $\measuredangle[q\,^p_y] = \pi$ and $\uparrow_{[qz]}$ is distinct from $\uparrow_{[qp]}$ and $\uparrow_{[qp]}$. Thus, according to Proposition 15.7, $\operatorname{LinDim}\mathcal{L} > 1$, a contradiction. $\triangle$

Clearly ❷ + ❸ ⇒ ❶; hence the result. $\square$

Chapter 16

Gradient flow

Gradient flow could be considered as a nonsmooth version of first-order ordinary differential equations. It provides a universal tool in Alexandrov geometry.

We consider only the CBB case since the main applications of gradient flow are there. But the proofs in this chapter admit straightforward extension to locally compact length spaces with defined angles between geodesics, as well as CAT spaces.

The technique of gradient flow takes its roots in Sharafutdinov's retraction, introduced by Vladimir Sharafutdinov [122]. It has been used widely in comparison geometry since then. In CBB spaces, it was first used by Grigory Perelman and the third author [101]. A bit later, independently Jürgen Jost and Uwe Mayer [72, 93] used the gradient flow in CAT spaces. Later, Alexander Lytchak unified and generalized these two approaches to a wide class of metric spaces in [90]. It was developed yet further by Shin-ichi Ohta [99] and by Giuseppe Sevaré [121].

The following exercise is a stripped-down version of Sharfutdinov's retraction; it gives the idea behind gradient flow.

16.1. Exercise. *Assume that a one-parameter family of convex sets $K_t \subset \mathbb{E}^m$ is nested; that is, $K_{t_1} \supset K_{t_2}$ if $t_1 \leqslant t_2$. Show that there is a family of short maps $\varphi_t \colon \mathbb{E}^m \to K_t$ such that $\varphi_t|_{K_t} = \mathrm{id}$ and $\varphi_{t_2} \circ \circ \varphi_{t_1} = \varphi_{t_2}$ for all $t_1 \leqslant t_2$.*

A Gradient-like curves

Gradient-like curves will be used later in the construction of gradient curves. The latter are a special reparametrization of gradient-like curves.

16.2. Definition. *Let $\mathcal{Z}$ be a complete length space and $f \colon \mathcal{Z} \hookrightarrow \mathbb{R}$ be* *locally Lipschitz semiconcave subfunction. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT.

A Lipschitz curve $\hat\alpha\colon [s_{\min}, s_{\max}) \to \operatorname{Dom} f$ will be called an f-gradient-like curve if

$$\hat\alpha^+ = \tfrac{1}{|\nabla_{\hat\alpha} f|}\cdot\nabla_{\hat\alpha} f;$$

that is, for any $s \in [s_{\min}, s_{\max})$, $\hat\alpha^+(s)$ is defined and

$$\hat\alpha^+(s) = \tfrac{1}{|\nabla_{\hat\alpha(s)} f|}\cdot\nabla_{\hat\alpha(s)} f.$$

Note that this definition implies that $|\nabla_p f| > 0$ for any point p on $\hat\alpha$.

The following theorem gives a seemingly weaker condition that is equivalent to the definition of gradient-like curve.

16.3. Theorem. *Suppose $\mathcal{Z}$ is a complete length space, $f\colon \mathcal{Z} \hookrightarrow \mathbb{R}$ is a locally Lipschitz semiconcave subfunction, and $|\nabla_p f| > 0$ for any $p \in$ $\in \operatorname{Dom} f$. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT.

A curve $\hat\alpha\colon [s_{\min}, s_{\max}) \to \operatorname{Dom} f$ is an f-gradient-like curve if and only if it is 1-Lipschitz and

$$\text{❶}\qquad \varliminf_{s\to s_0+} \frac{f\circ\hat\alpha(s) - f\circ\hat\alpha(s_0)}{s - s_0} \geqslant |\nabla_{\hat\alpha(s_0)} f|$$

for almost all $s_0 \in [s_{\min}, s_{\max})$.

Proof. The "only if" part follows directly from the definition. To prove the "if" part, note that for any $s_0 \in [s_{\min}, s_{\max})$ we have

$$\begin{aligned}\varliminf_{s\to s_0+} \frac{f\circ\hat\alpha(s) - f\circ\hat\alpha(s_0)}{s - s_0} &\geqslant \varliminf_{s\to s_0+} ⨍_{s_0}^{s} |\nabla_{\hat\alpha(\varsigma)} f|\cdot d\varsigma \geqslant \\ &\geqslant |\nabla_{\hat\alpha(s_0)} f|;\end{aligned}$$

the first inequality follows from ❶ and the second from lower semicontinuity of the function $x \mapsto |\nabla_x f|$, see 13.29. From 13.20, we have

$$\hat\alpha^+(s_0) = \tfrac{1}{|\nabla_{\hat\alpha(s_0)} f|}\cdot\nabla_{\hat\alpha(s_0)} f.$$

Hence the result. □

Recall that second-order differential inequalities are understood in a barrier sense; see Section 5E.

16.4. Theorem. *Let $\mathcal{Z}$ be a complete length space and $f\colon \mathcal{Z} \hookrightarrow \mathbb{R}$ be locally Lipschitz and λ-concave. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT. *Assume $\hat\alpha\colon [0, s_{\max}) \to \operatorname{Dom} f$ is an f-gradient-like curve. Then*

$$(f\circ\hat\alpha)'' \leqslant \lambda$$

everywhere on $[0, s_{\max})$.

Closely related statements were proved independently by Uwe Mayer and Shin-ichi Ohta [93, 2.36] and [99, 5.7].

Before the proof, let us formulate and prove a corollary.

16.5. Corollary. *Let $\mathcal{Z}$ be a complete length space, $f\colon \mathcal{Z} \hookrightarrow \mathbb{R}$ be a locally Lipschitz and semiconcave function, and $\hat\alpha\colon [0, s_{\max}) \to \operatorname{Dom} f$ be an f-gradient-like curve. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT. *Then the function $s \mapsto |\nabla_{\hat\alpha(s)} f|$ is right-continuous; that is, for any $s_0 \in [0, s_{\max})$ we have*

$$|\nabla_{\hat\alpha(s_0)} f| = \lim_{s \to s_0+} |\nabla_{\hat\alpha(s)} f|.$$

Proof. Applying 16.4 locally, we have that $f \circ \hat\alpha(s)$ is semiconcave. The statement follows since

$$(f \circ \hat\alpha)^+(s) = (\boldsymbol{d}_p f)\left(\tfrac{1}{|\nabla_{\hat\alpha(s)} f|}\cdot\nabla_{\hat\alpha(s)} f\right) = |\nabla_{\hat\alpha(s)} f|.$$

□

Proof of 16.4. For any $s > s_0$,

$$\begin{aligned}(f \circ \hat\alpha)^+(s_0) = |\nabla_{\hat\alpha(s_0)} f| &\geqslant \\ &\geqslant (d_{\hat\alpha(s_0)} f)(\uparrow_{[\hat\alpha(s_0)\hat\alpha(s)]}) \geqslant \\ &\geqslant \frac{f\circ\hat\alpha(s) - f\circ\hat\alpha(s_0)}{|\hat\alpha(s) - \hat\alpha(s_0)|} - \tfrac{\lambda}{2}\cdot|\hat\alpha(s) - \hat\alpha(s_0)|.\end{aligned}$$

Let $\lambda_+ = \max\{0, \lambda\}$. Since $s - s_0 \geqslant |\hat\alpha(s) - \hat\alpha(s_0)|$, for any $s > s_0$ we have

❷ $$(f \circ \hat\alpha)^+(s_0) \geqslant \frac{f\circ\hat\alpha(s) - f\circ\hat\alpha(s_0)}{s - s_0} - \tfrac{\lambda_+}{2}\cdot(s - s_0).$$

Thus $f \circ \hat\alpha$ is λ_+-concave. That finishes the proof for $\lambda \geqslant 0$. For $\lambda < 0$ we get only that $f \circ \hat\alpha$ is 0-concave.

Note that $|\hat\alpha(s) - \hat\alpha(s_0)| = s - s_0 - o(s - s_0)$. Thus

❸ $$(f \circ \hat\alpha)^+(s_0) \geqslant \frac{f\circ\hat\alpha(s) - f\circ\hat\alpha(s_0)}{s - s_0} - \tfrac{\lambda}{2}\cdot(s - s_0) + o(s - s_0).$$

Together, ❷ and ❸ imply that $f \circ \hat\alpha$ is λ-concave. □

16.6. Proposition. *Let $\mathcal{L}$ be a complete length* CBB(κ) *space, $p, q \in \in \mathcal{L}$. Assume $\hat\alpha\colon [s_{\min}, s_{\max}) \to \mathcal{L}$ is a dist_p-gradient-like curve such that $\hat\alpha(s) \to z \in \,]pq[$ as $s \to s_{\max}+$. Then $\hat\alpha$ is a unit-speed geodesic that lies in $[pq]$.*

Proof. Clearly,

$$\textbf{❹} \qquad \tfrac{d^+}{dt}|q-\hat\alpha(t)| \geqslant -1.$$

On the other hand,

$$\textbf{❺} \qquad \begin{aligned}\tfrac{d^+}{dt}|p-\hat\alpha(t)| &\geqslant (\boldsymbol{d}_{\hat\alpha(t)}\operatorname{dist}_p)(\uparrow_{[\hat\alpha(t)q]}) \geqslant \\ &\geqslant -\cos\tilde\measuredangle^{\kappa}(\hat\alpha(t)\,{}^{p}_{q}).\end{aligned}$$

Inequalities ❹ and ❺ imply that the function $t \mapsto \tilde\measuredangle^{\kappa}(q\,{}^{\hat\alpha(t)}_{p})$ is nondecreasing. Hence the result. □

B Gradient curves

In this section we define gradient curves and tie them tightly to gradient-like curves which were introduced in Section 16A.

16.7. Definition. *Let $\mathcal{Z}$ be a complete length space and $f\colon \mathcal{Z} \hookrightarrow \mathbb{R}$ be a locally Lipschitz and semiconcave function. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT.

A locally Lipschitz curve $\alpha\colon [t_{\min}, t_{\max}) \to \operatorname{Dom} f$ will be called an f-gradient curve if

$$\alpha^+ = \nabla_\alpha f;$$

that is, for any $t \in [t_{\min}, t_{\max})$, $\alpha^+(t)$ is defined and $\alpha^+(t) = \nabla_{\alpha(t)} f$.

The following exercise describes a global geometric property of a gradient curve without direct reference to its function. It uses the notion of self-contracting curves which was introduced by Aris Daniilidis, Olivier Ley, Stéphane Sabourau [44].

16.8. Exercise. *Let $\mathcal{Z}$ be a complete length space, $f\colon \mathcal{Z} \to \mathbb{R}$ a concave locally Lipschitz function, and $\alpha\colon \mathbb{I} \to \mathcal{Z}$ an f-gradient curve. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT.

Show that α is self-contracting; that is,

$$|\alpha(t_1) - \alpha(t_3)|_{\mathcal{Z}} \geqslant |\alpha(t_2) - \alpha(t_3)|_{\mathcal{Z}}$$

if $t_1 \leqslant t_2 \leqslant t_3$.

The next lemma states that gradient and gradient-like curves are special reparametrizations of each other.

16.9. Lemma. *Let $\mathcal{Z}$ be a complete length space and $f\colon \mathcal{Z} \hookrightarrow \mathbb{R}$ be a locally Lipschitz semiconcave subfunction such that $|\nabla_p f| > 0$ for any $p \in \operatorname{Dom} f$. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT.

Assume that $\alpha\colon [0, t_{\max}) \to \operatorname{Dom} f$ *is a locally Lipschitz curve and* $\hat\alpha\colon [0, s_{\max}) \to \operatorname{Dom} f$ *is its reparametrization by arc-length, so* $\alpha = \hat\alpha \circ \sigma$ *for some homeomorphism* $\sigma\colon [0, t_{\max}) \to [0, s_{\max})$. *Then*

$$\alpha^+ = \nabla_\alpha f$$
$$\Updownarrow$$
$$\hat\alpha^+ = \frac{1}{|\nabla_{\hat\alpha} f|}\cdot\nabla_{\hat\alpha} f \quad \textit{and} \quad \sigma^{-1}(s) = \int\limits_0^s \frac{\boldsymbol{d}ş}{(f\circ\hat\alpha)'(ş)}.$$

Proof; $(\Rightarrow)$. According to 5.10,

❶
$$\begin{aligned}\sigma'(t) &\overset{a.e.}{=\!=} |\alpha^+(t)| = \\ &= |\nabla_{\alpha(t)} f|.\end{aligned}$$

Note that

$$\begin{aligned}(f\circ\alpha)'(t) &\overset{a.e.}{=\!=} (f\circ\alpha)^+(t) = \\ &= |\nabla_{\alpha(t)} f|^2.\end{aligned}$$

Setting $s = \sigma(t)$, we have

$$\begin{aligned}(f\circ\hat\alpha)'(s) &\overset{a.e.}{=\!=} \frac{(f\circ\alpha)'(t)}{\sigma'(t)} \overset{a.e.}{=\!=} \\ &\overset{a.e.}{=\!=} |\nabla_{\alpha(t)} f| = \\ &= |\nabla_{\hat\alpha(s)} f|.\end{aligned}$$

From 16.3, it follows that $\hat\alpha(t)$ is an f-gradient-like curve; that is,

$$\hat\alpha^+ = \frac{1}{|\nabla_{\hat\alpha} f|}\cdot\nabla_{\hat\alpha} f.$$

In particular, $(f\circ\hat\alpha)^+(s) = |\nabla_{\hat\alpha^+(s)} f|$, and by ❶,

$$\begin{aligned}\sigma^{-1}(s) &= \int\limits_0^s \frac{1}{|\nabla_{\hat\alpha(ş)} f|}\cdot\boldsymbol{d}ş = \\ &= \int\limits_0^s \frac{1}{(f\circ\hat\alpha)'(ş)}\cdot\boldsymbol{d}ş.\end{aligned}$$

$(\Leftarrow)$. Clearly,

$$\begin{aligned}\sigma(t) &= \int\limits_0^t (f\circ\hat\alpha)^+(\sigma(ţ))\cdot\boldsymbol{d}ţ = \\ &= \int\limits_0^t |\nabla_{\alpha(ţ)} f|\cdot\boldsymbol{d}ţ.\end{aligned}$$

According to 16.5, the function $s \mapsto |\nabla_{\hat\alpha(s)} f|$ is right-continuous. Therefore so is the function $t \mapsto |\nabla_{\hat\alpha\circ\sigma(t)} f| = |\nabla_{\alpha(t)} f|$. Hence, for any $t_0 \in \in [0, t_{\max})$ we have

$$\begin{aligned}\sigma^+(t_0) &= \lim_{t\to t_0+} \fint_{t_0}^{t} |\nabla_{\alpha(\tau)} f|\cdot d\tau = \\ &= |\nabla_{\alpha(t_0)} f|.\end{aligned}$$

Thus, we have

$$\begin{aligned}\alpha^+(t_0) &= \sigma^+(t_0)\cdot\hat\alpha^+(\sigma(t_0)) = \\ &= \nabla_{\alpha(t_0)} f.\end{aligned}$$

□

16.10. Exercise. *Let $\mathcal{Z}$ be a complete length space, and $f\colon \mathcal{Z} \to \mathbb{R}$ be a semiconcave locally Lipschitz function. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT. *Assume $\alpha\colon \mathbb{I} \to \mathcal{Z}$ is a Lipschitz curve such that*

$$\begin{aligned}\alpha^+(t) &\leqslant |\nabla_{\alpha(t)} f|,\\ (f\circ\alpha)^+(t) &\geqslant |\nabla_{\alpha(t)} f|^2\end{aligned}$$

for almost all t. Show that α is an f-gradient curve.

16.11. Exercise. *Let $\mathcal{Z}$ be a complete length space and $f\colon \mathcal{Z} \to \mathbb{R}$ be a concave locally Lipschitz function. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT. *Show that $\alpha\colon \mathbb{R} \to \mathcal{Z}$ is an f-gradient curve if and only if*

$$|x - \alpha(t_1)|^2_{\mathcal{Z}} - |x - \alpha(t_0)|^2_{\mathcal{Z}} \leqslant 2\cdot(t_1 - t_0)\cdot(f\circ\alpha(t_1) - f(x))$$

for any $t_1 > t_0$ and $x \in \mathcal{Z}$.

Distance estimates

16.12. First distance estimate. *Let $\mathcal{Z}$ be a complete length space, and $f\colon \mathcal{Z} \to \mathbb{R}$ be a locally Lipschitz λ-concave function. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT. *Let $\alpha, \beta\colon [0, t_{\max}) \to \mathcal{Z}$ be two f-gradient curves. Then*

$$|\alpha(t) - \beta(t)| \leqslant e^{\lambda\cdot t}\cdot|\alpha(0) - \beta(0)|$$

for any t.

Moreover, the statement holds for a locally Lipschitz λ-concave subfunction $f\colon \mathcal{Z} \hookrightarrow \mathbb{R}$ if there is a geodesic $[\alpha(t)\,\beta(t)]$ in Dom *f for any t.*

Proof. If $\mathcal{Z}$ is not geodesic, then pass to its ultrapower $\mathcal{Z}^{\omega}$.

Fix a choice of geodesic $[\alpha(t)\,\beta(t)]$ for each t.

Setting $\ell(t) = |\alpha(t) - \beta(t)|$, from the first variation inequality (6.6) and the estimate in 13.23 we get

$$\ell^+(t) \leqslant -\langle\uparrow_{[\alpha(t)\beta(t)]}, \nabla_{\alpha(t)}f\rangle - \langle\uparrow_{[\beta(t)\alpha(t)]}, \nabla_{\beta(t)}f\rangle \leqslant \lambda{\cdot}\ell(t).$$

Here one has to choose a midpont p of $[\alpha(t)\,\beta(t)]$, apply the first variation inequality for distance to p, and apply the triangle inequality. Hence the result. □

16.13. Second distance estimate. *Let $\mathcal{Z}$ be a complete length space, $\varepsilon > 0$, and $f, g\colon \mathcal{Z} \to \mathbb{R}$ be two λ-concave locally Lipschitz function such that $|f - g| < \varepsilon$. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT. *Assume $\alpha, \beta\colon [0, t_{\max}) \to \mathcal{Z}$ are respectively f- and g-gradient curves such that $\alpha(0) = \beta(0)$. Then*

$$|\alpha(t) - \beta(t)| \leqslant \sqrt{\tfrac{1}{2\cdot\varepsilon\cdot\lambda}\cdot\left(e^{\frac{t\cdot\lambda}{\varepsilon}} - 1\right)}$$

for any $t \in [0, t_{\max})$. In particular, if $t_{\max} < \infty$ then

$$|\alpha(t) - \beta(t)| \leqslant c{\cdot}\sqrt{\varepsilon{\cdot}t}$$

for some constant $c = c(t_{\max}, \lambda)$.

Moreover, the same conclusion holds for locally Lipschitz λ-concave subfunctions $f, g\colon \mathcal{Z} \hookrightarrow \mathbb{R}$ if for any $t \in [0, t_{\max})$ there is a geodesic $[\alpha(t)\,\beta(t)]$ in $\operatorname{Dom} f \cap \operatorname{Dom} g$.

Proof. Set $\ell = \ell(t) = |\alpha(t) - \beta(t)|$. Fix t, and let $p = \alpha(t)$ and $q = \beta(t)$. From the first variation formula and 13.22,

$$\begin{aligned}\ell^+ &\leqslant -\langle\uparrow_{[pq]}, \nabla_p f\rangle - \langle\uparrow_{[qp]}, \nabla_q g\rangle \leqslant \\ &\leqslant -\left(f(q) - f(p) - \lambda{\cdot}\tfrac{\ell^2}{2}\right)/\ell - \left(g(p) - g(q) - \lambda{\cdot}\tfrac{\ell^2}{2}\right)/\ell \leqslant \\ &\leqslant \lambda{\cdot}\ell + \tfrac{2\cdot\varepsilon}{\ell}.\end{aligned}$$

Integrating the above estimate, we get

$$\ell(t) \leqslant \sqrt{\tfrac{1}{2\cdot\varepsilon\cdot\lambda}\cdot\left(e^{\frac{t\cdot\lambda}{\varepsilon}} - 1\right)}.$$

□

Existence, uniqueness, convergence

In general, the "past" of gradient curves can not be determined by the present. For example, consider the concave function $f\colon \mathbb{R}\to\mathbb{R}$, $f(x) = -|x|$. The two curves $\alpha(t)=\min\{0,t\}$ with $\beta(t)=0$ are f-gradient with $\alpha(t)=\beta(t)=0$ for all $t\geqslant 0$; however $\alpha(t)\neq\beta(t)$ for all $t<0$.

The next theorem shows that a "future" gradient curve is unique.

16.14. Picard's theorem. *Let $\mathcal{Z}$ be a complete length space and $f\colon \mathcal{Z}\hookrightarrow\mathbb{R}$ be a semiconcave subfunction. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT. *Assume $\alpha,\beta\colon [0,t_{\max})\to \operatorname{Dom} f$ are two f-gradient curves such that $\alpha(0)=\beta(0)$. Then $\alpha(t)=\beta(t)$ for any $t\in[0,t_{\max})$.*

Proof. This follows from the first distance estimate (16.12). □

16.15. Local existence. *Let $\mathcal{Z}$ be a complete length space and $f\colon \mathcal{Z}\hookrightarrow\mathbb{R}$ be locally Lipschitz λ-concave subfunction. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT. *Then for any $p\in\operatorname{Dom} f$,*

a) if $|\nabla_p f|>0$, then for some $\varepsilon>0$, there is an f-gradient-like curve $\hat\alpha\colon [0,\varepsilon)\to\mathcal{Z}$ that starts at p (that is, $\hat\alpha(0)=p$);

b) for some $\delta>0$, there is an f-gradient curve $\alpha\colon[0,\delta)\to\mathcal{Z}$ that starts at p (that is $\alpha(0)=p$).

This theorem was proved by Grigory Perelman and the third author [101]; we present a simplified proof given by Alexander Lytchak [90].

Proof. If $|\nabla_p f|=0$, then the constant curve $\alpha(t)=p$ is f-gradient.

Otherwise, choose $\varepsilon>0$ such that $\mathrm{B}(p,\varepsilon)\subset\operatorname{Dom} f$, the restriction $f|_{\mathrm{B}(p,\varepsilon)}$ is Lipschitz, and $|\nabla_x f|>\varepsilon$ for all $x\in\mathrm{B}(p,\varepsilon)$; the latter is possible due to semicontinuity of |gradient| (13.29).

The curves $\hat\alpha$ and α will be constructed in the following three steps. First we construct an f^ω-gradient-like curve $\hat\alpha_\omega\colon[0,\varepsilon)\to\mathcal{Z}^\omega$ as an ω-limit of a certain sequence of broken geodesics in $\mathcal{Z}$. Second, we parametrize $\hat\alpha_\omega$ as in 16.9, to obtain an f^ω-gradient curve α_ω in $\mathcal{Z}^\omega$. Third, applying Picard's theorem (16.14) together with Lemma 3.4, we obtain that α_ω lies in $\mathcal{Z}\subset\mathcal{Z}^\omega$ and therefore one can take $\alpha=\alpha_\omega$ and $\hat\alpha=\hat\alpha_\omega$.

Note that if $\mathcal{Z}$ is proper, then $\mathcal{Z}=\mathcal{Z}^\omega$ and $f^\omega=f$. Thus, in this case, the third step is not necessary.

Step 1. Given $n\in\mathbb{N}$, by an open-closed argument, we can construct a unit-speed curve $\hat\alpha_n\colon[0,\varepsilon]\to\mathcal{Z}$ starting at p, with a partition of $[0,\varepsilon)$ into a countable number of half-open intervals $[\sigma_i,\bar\sigma_i)$ such that for each i we have

(i) $\hat\alpha_n([\sigma_i,\bar\sigma_i])$ is a geodesic and $\bar\sigma_i-\sigma_i<\frac{1}{n}$,

(ii) $f\circ\hat\alpha_n(\bar\sigma_i)-f\circ\hat\alpha_n(\sigma_i)>(\bar\sigma_i-\sigma_i)\cdot(|\nabla_{\hat\alpha_n(\sigma_i)}f|-\frac{1}{n})$.

Passing to a subsequence of $\hat{\alpha}_n$ such that $f \circ \hat{\alpha}_n$ uniformly converges, let

$$h(s) = \lim_{n \to \infty} f \circ \hat{\alpha}_n(s).$$

Let $\hat{\alpha}_\omega = \lim_{n \to \omega} \hat{\alpha}_n$, a curve in $\mathcal{Z}^\omega$ that starts at $p \in \mathcal{Z} \subset \mathcal{Z}^\omega$. Clearly $\hat{\alpha}_\omega$ is 1-Lipschitz. From (ii) and 13.26, we have

$$(f^\omega \circ \hat{\alpha}_\omega)^+(\sigma) \geqslant |\nabla_{\hat{\alpha}_\omega(\sigma)} f^\omega|.$$

According to 16.3, $\hat{\alpha}_\omega \colon [0, \varepsilon) \to \mathcal{Z}^\omega$ is an f^ω-gradient-like curve.

Step 2. Clearly $h(s) = f^\omega \circ \alpha_\omega$. Therefore, according to 16.4, h is λ-concave. Thus we can define a homeomorphism $\sigma \colon [0, \delta] \to [0, \varepsilon]$ by

❷ $$\sigma^{-1}(s) = \int_0^s \frac{1}{h'(\varsigma)} \cdot \boldsymbol{d}\varsigma,$$

According to 16.9, $\alpha(t) = \hat{\alpha} \circ \sigma(t)$ is an f^ω-gradient curve in $\mathcal{Z}^\omega$.

Step 3. Clearly, $\nabla_p f = \nabla_p f^\omega$ for any $p \in \mathcal{Z} \subset \mathcal{Z}^\omega$; more formally, if $\iota \colon \mathcal{Z} \hookrightarrow \mathcal{Z}^\omega$ is the natural embedding, then $(\boldsymbol{d}_p \iota)(\nabla_p f) = \nabla_p f^\omega$. Thus it is sufficient to show that α_ω lies in $\mathcal{Z}$. Assume the contrary; then according to 3.4, there is a subsequence $\hat{\alpha}_{n_k}$ such that

$$\hat{\alpha}_\omega \neq \hat{\alpha}'_\omega := \lim_{k \to \omega} \hat{\alpha}_{n_k}.$$

Clearly $h(s) = f^\omega \circ \hat{\alpha}_\omega = f^\omega \circ \hat{\alpha}'_\omega$. Thus for $\sigma \colon [0, \delta] \to [0, \varepsilon]$ defined by ❷, we have that both curves $\hat{\alpha}_\omega \circ \sigma$ and $\hat{\alpha}'_\omega \circ \sigma$ are f^ω-gradient. From Picard's theorem (16.14), we have $\hat{\alpha}_\omega \circ \sigma = \hat{\alpha}'_\omega \circ \sigma$. Therefore $\hat{\alpha}_\omega = \hat{\alpha}'_\omega$, a contradiction. □

16.16. Ultralimit of gradient curves. *Assume*

- ◇ $\mathcal{Z}_n$ *is a sequence of complete spaces,* $\mathcal{Z}_n \to \mathcal{Z}_\omega$ *as* $n \to \omega$, *and* $p_n \to p_\omega$ *for a sequence of points* $p_n \in \mathcal{Z}_n$,
- ◇ *all spaces* $\mathcal{Z}_n$ *are either* CBB(κ) *or* CAT(κ),
- ◇ $f_n \colon \mathcal{Z}_n \looparrowright \mathbb{R}$ *are* ℓ*-Lipschitz and* λ*-concave,* $f_n \to f_\omega$ *as* $n \to \omega$, *and* $p_\omega \in \operatorname{Dom} f_\omega$.

Then:

a) f_ω *is* λ*-concave.*

b) If $|\nabla_{p_\omega} f_\omega| > 0$, *then there is* $\varepsilon > 0$ *such that, the* f_n*-gradient-like curves* $\hat{\alpha}_n \colon [0, \varepsilon) \to \mathcal{Z}_n$ *are defined for* ω*-almost all* n. *Moreover, a curve* $\hat{\alpha}_\omega \colon [0, \varepsilon) \to \mathcal{Z}_\omega$ *is a gradient-like curve that starts at* p_ω *if and only if* $\hat{\alpha}_n(s) \to \hat{\alpha}_\omega(s)$ *as* $n \to \omega$ *for all* $s \in [0, \varepsilon)$.

c) For some $\delta > 0$, the f_n-gradient curves $\alpha_n\colon [0,\delta) \to \mathcal{Z}_n$ are defined for ω-almost all n. Moreover, a curve $\alpha_\omega\colon [0,\delta) \to \mathcal{Z}_\omega$ is a gradient curve that starts at p_ω if and only if $\alpha_n(t) \to \to \alpha_\omega(t)$ as $n \to \omega$ for all $t \in [0,\delta)$.

Note that according to Exercise 3.10, part (a) does not hold for general metric spaces. The idea of the proof is the same as in the proof of local existence (16.15).

Proof of 16.16; (a) Fix a geodesic $\gamma_\omega\colon \mathbb{I} \to \operatorname{Dom} f_\omega$; we need to show that the function

$$t \mapsto f_\omega \circ \gamma_\omega(t) - \tfrac{\lambda}{2}\cdot t^2 \qquad ❸$$

is concave.

Since the f_n are ℓ-Lipschitz, so is f_ω. Therefore it is sufficient to prove concavity in the interior of $\mathbb{I}$. In particular we can assume that γ_ω is sufficiently short and can be extended behind its ends p_ω and q_ω as a minimizing geodesic. If $\mathcal{Z}$ is CBB, then by Theorem 8.11, γ_ω is the unique geodesic connecting p_ω to q_ω. The same holds true if $\mathcal{Z}$ is CAT by the uniqueness of geodesics (9.8).

Construct two sequences of points $p_n, q_n \in \mathcal{Z}_n$ such that $p_n \to p_\omega$ and $q_n \to q_\omega$ as $n \to \omega$. Applying either 8.11 or 9.8, we can assume that for each n there is a geodesic γ_n from p_n to q_n in $\mathcal{Z}_n$.

Since f_n is λ-concave, the function

$$t \mapsto f_n \circ \gamma_n(t) - \tfrac{\lambda}{2}\cdot t^2$$

is concave.

The ω-limit of the sequence γ_n is a geodesic in $\mathcal{Z}_\omega$ from p_ω to q_ω. By uniqueness of such geodesics, we have that $\gamma_n \to \gamma_\omega$ as $n \to \omega$. Passing to the limit, we have ❸.

"if"-part of (b). Take $\varepsilon > 0$ so small that $\mathrm{B}(p_\omega, \varepsilon) \subset \operatorname{Dom} f_\omega$ and $|\nabla_{x_\omega} f_\omega| > 0$ for any $x_\omega \in \mathrm{B}(p_\omega, \varepsilon)$ (this is possible by 13.29).

Clearly $\hat\alpha_\omega$ is 1-Lipschitz. From 13.26, we get $(f_\omega \circ \hat\alpha_\omega)^+(s) \geqslant \geqslant |\nabla_{\hat\alpha_\omega(s)} f^\omega|$. According to 16.3, $\hat\alpha_\omega\colon [0,\varepsilon) \to \mathcal{Z}^\omega$ is an f_ω-gradient-like curve.

"if"-part of (c). Assume first that $|\nabla_{p_\omega} f_\omega| > 0$, so we can apply the "if"-part of (b). Let $h_n = f_n \circ \hat\alpha_n\colon [0,\varepsilon) \to \mathbb{R}$ and $h_\omega = f_\omega \circ \hat\alpha_\omega$. From 16.4, the h_n are λ-concave, and clearly $h_n \to h_\omega$ as $n \to \omega$. Let us define reparametrizations

$$\sigma_n^{-1}(s) = \int_0^s \frac{1}{h_n'(\varsigma)}\cdot \boldsymbol{d}\varsigma, \qquad \sigma_\omega^{-1}(s) = \int_0^s \frac{1}{h_\omega'(\varsigma)}\cdot \boldsymbol{d}\varsigma.$$

The λ-convexity of the h_n implies that $\sigma_n \to \sigma_\omega$ as $n \to \omega$. By 16.9, $\alpha_n = \hat{\alpha}_n \circ \sigma_n$. Applying the "if"-part of (b) together with Lemma 16.9, we get that $\alpha_\omega = \hat{\alpha}_\omega \circ \sigma_\omega$ is gradient curve.

The remaining case $|\nabla_{p_\omega} f_\omega| = 0$ can be reduced to the one above using the following trick. Consider the sequence of spaces $\mathcal{Z}_n^\times = \mathcal{Z}_n \times \mathbb{R}$, with the sequence of subfunctions $f_n^\times : \mathcal{Z}_n^\times \to \mathbb{R}$ defined by

$$f_n^\times(p,t) = f_n(p) + t.$$

Applying either 11.7b or 11.6b, we have that $\mathcal{Z}_n^\times$ is a CBB(κ_-) space for $\kappa_- = \min\{\kappa, 0\}$, or CAT($\kappa_+$) space for $\kappa_+ = \max\{\kappa, 0\}$. Note that the $f_n^\times$ are λ_+-concave for $\lambda_+ = \max\{\lambda, 0\}$. Now let $\mathcal{Z}_\omega^\times = \mathcal{Z}_\omega \times \mathbb{R}$, and $f_\omega^\times(p,t) = f_\omega(p) + t$.

Clearly $\mathcal{Z}_n^\times \to \mathcal{Z}_\omega^\times$, $f_n^\times \to f_\omega^\times$ as $n \to \omega$, and $|\nabla_x f_\omega^\times| > 0$ for any $x \in \operatorname{Dom} f_\omega^\times$. Thus for the sequence $f_n^\times : \mathcal{Z}_n^\times \multimap \mathbb{R}$, we can apply the "if"-part of (b). It remains to note that the curve $\alpha_\omega^\times(t) = (\alpha_\omega(t), t)$ is an $f_\omega^\times$-gradient curve in $\mathcal{Z}_\omega^\times$ if and only if $\alpha_\omega(t)$ is an f_ω-gradient curve.

"only if"-part of (c) and (b). The "only if"-part of (c) follows from the "if"-part of (c) and Picard's theorem (16.14). Applying Lemma 16.9, we get the "only if"-part of (b). □

Directly from local existence (16.15) and the distance estimates (16.12), we obtain:

16.17. Global existence. *Let $f\colon \mathcal{Z} \multimap \mathbb{R}$ be a locally Lipschitz and λ-concave subfunction on a complete length space $\mathcal{Z}$. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT*. Then for any $p \in \operatorname{Dom} f$, there is $t_{\max} \in (0, \infty]$ such that there is an f-gradient curve $\alpha\colon [0, t_{\max}) \to \mathcal{Z}$ with $\alpha(0) = p$. Moreover, for any sequence $t_n \to t_{\max}-$, the sequence $\alpha(t_n)$ does not have a limit point in* Dom f.

The following theorem guarantees the existence of gradient curves for all times for the special type of semiconcave functions that play important role in the theory. It follows from 16.17, 16.4 and 16.9.

16.18. Theorem. *Let $\mathcal{Z}$ be a complete length space and $f\colon \mathcal{Z} \to \mathbb{R}$ satisfies*

$$f'' + \kappa{\cdot}f \leqslant \lambda$$

for some real values κ and λ. Suppose that $\mathcal{Z}$ is either CBB *or* CAT*. Then f has* c o m p l e t e g r a d i e n t*; that is, for any $x \in \mathcal{Z}$ there is a f-gradient curve $\alpha\colon [0, \infty) \to \mathcal{Z}$ that starts at x.*

C Gradient flow

In this section we define gradient flow for semiconcave subfunctions and reformulate theorems obtained earlier in this chapter using this new terminology.

Let $\mathcal{Z}$ be a complete length space and $f\colon \mathcal{Z} \hookrightarrow \mathbb{R}$ be a locally Lipschitz semiconcave subfunction. Suppose that $\mathcal{Z}$ is either CBB or CAT. For any $t \geqslant 0$, we write $\mathrm{Flow}_f^t(x) = y$ if there is an f-gradient curve α such that $\alpha(0) = x$ and $\alpha(t) = y$. The partially defined map Flow_f^t from $\mathcal{Z}$ to itself is called the f-gradient flow for time t.

From 16.13, it follows that for any $t \geqslant 0$, the domain of definition of Flow_f^t is an open subset of $\mathcal{Z}$; that is, Flow_f^t is a submap. Moreover, if f is defined on all of $\mathcal{Z}$ and $f'' + \mathrm{K}\cdot f \leqslant \lambda$ for some $\mathrm{K}, \lambda \in \mathbb{R}$, then according to 16.18, $\mathrm{Flow}_f^t(x)$ is defined for all pairs $(x,t) \in \mathcal{Z} \times \mathbb{R}_{\geqslant 0}$.

Clearly $\mathrm{Flow}_f^{t_1+t_2} = \mathrm{Flow}_f^{t_1} \circ \mathrm{Flow}_f^{t_2}$; in other words, gradient flow is given by an action of the semigroup $(\mathbb{R}_{\geqslant 0}, +)$.

From the first distance estimate (16.12), we have the following:

16.19. Proposition. *Let $\mathcal{Z}$ be a complete length* CBB *space and $f\colon \mathcal{Z} \to \mathbb{R}$ be a semiconcave function. Then the map $x \mapsto \mathrm{Flow}_f^t(x)$ is locally Lipschitz.*

Moreover, if f is λ-concave, then Flow_f^t is $e^{\lambda\cdot t}$-Lipschitz.

The next proposition states that gradient flow is stable under Gromov–Hausdorff convergence. The proposition follows directly from the proposition on ultralimit of gradient curves 16.16.

16.20. Proposition. *If $\mathcal{Z}_n$ is an m-dimensional complete length* $\mathrm{CBB}(\kappa)$ *space, $\mathcal{Z}_n \xrightarrow{\tau} \mathcal{Z}$, and $f_n\colon \mathcal{Z}_n \to \mathbb{R}$ is a sequence of λ-concave functions that converges to $f\colon \mathcal{Z} \to \mathbb{R}$, then $\mathrm{Flow}_{f_n}^t\colon \mathcal{Z}_n \to \mathcal{Z}_n$ converges to $\mathrm{Flow}_f^t\colon \mathcal{Z} \to \mathcal{Z}$.*

D Line splitting theorem

Let $\mathcal{X}$ be a metric space and $A, B \subset \mathcal{X}$. We will write

$$\mathcal{X} = A \oplus B$$

if there exist projections $\mathrm{proj}_A\colon \mathcal{X} \to A$ and $\mathrm{proj}_B\colon \mathcal{X} \to B$ such that

$$|x-y|^2 = |\mathrm{proj}_A(x) - \mathrm{proj}_A(y)|^2 + |\mathrm{proj}_B(x) - \mathrm{proj}_B(y)|^2$$

for any two points $x, y \in \mathcal{X}$.

Note that if

$$\mathcal{X} = A \oplus B$$

then

- A intersects B at a single point,
- both sets A and B are convex sets in $\mathcal{X}$.

Recall that a line in a metric space is a both-sided infinite geodesic; thus it minimizes the length on each segment.

16.21. Line splitting theorem. *Let $\mathcal{L}$ be a complete length* CBB(0) *space and γ be a line in $\mathcal{L}$. Then*

$$\mathcal{L} = X \oplus \gamma(\mathbb{R})$$

for some subset $X \subset \mathcal{L}$.

For smooth 2-dimensional surfaces, this theorem was proved by Stefan Cohn-Vossen [43]. For Riemannian manifolds of higher dimensions it was proved by Victor Toponogov [125]. Then it was generalized by Anatoliy Milka [95] to Alexandrov spaces; almost the same proof is given in [27, 1.5].

Further generalizations of the splitting theorem for Riemannian manifolds with nonnegative Ricci curvature were obtained by Jeff Cheeger and Detlef Gromoll [42]. This was further generalized by Jeff Cheeger and Toby Colding for limits of Riemannian manifolds with almost nonnegative Ricci curvature [40] and to their synthetic generalizations, so-called RCD spaces, by Nicola Gigli [53, 54]. Jost-Hinrich Eschenburg obtained an analogous result for Lorentzian manifolds [50], that is, pseudo-Riemannian manifolds of signature $(1, n)$.

We present a proof that uses gradient flow for Busemann functions. It is close in spirit to the proof given in [42].

Before going into the proof, let us state a few corollaries of the theorem.

16.22. Corollary. *Let $\mathcal{L}$ be a complete length* CBB(0) *space. Then there is an isometric splitting*

$$\mathcal{L} = \mathcal{L}' \oplus H$$

where $H \subset \mathcal{L}$ is a subset isometric to a Hilbert space, and $\mathcal{L}' \subset \mathcal{L}$ is a convex subset that contains no line.

16.23. Corollary. *Let $\mathcal{K}$ be an n-dimensional complete length* CBB(0) *cone and $v_+, v_- \in \mathcal{K}$ be a pair of opposite vectors (that is, $v_+ + v_- = 0$, see Definiton 13.33). Then there is an isometry $\iota\colon K \to K' \times \mathbb{R}$, where K' is a complete length* CBB(0) *space having a cone structure with tip $0'$ such that $\iota : v_\pm \mapsto (0', \pm|v_\pm|)$.*

16.24. Corollary. *Assume $\mathcal{L}$ is an m-dimensional complete length* CBB(1) *space, $m \geqslant 2$, and $\operatorname{rad}\mathcal{L} = \pi$. Then*

$$\mathcal{L} \stackrel{iso}{=\!=} \mathbb{S}^m.$$

The following lemma is closely relevant to the first distance estimate (16.12); its proof goes along the same lines.

16.25. Lemma. *Suppose $f\colon\mathcal{L}\to\mathbb{R}$ be a concave 1-Lipschitz function. Consider two f-gradient curves α and β. Then for any $t,s\geqslant 0$ we have*

$$|\alpha(s)-\beta(t)|^2\leqslant|p-q|^2+2\cdot(f(p)-f(q))\cdot(s-t)+(s-t)^2,$$

where $p=\alpha(0)$ and $q=\beta(0)$.

Proof. If $\mathcal{L}$ is not geodesic, then pass to its ultrapower $\mathcal{L}^{\omega}$.

Since f is 1-Lipschitz, $|\nabla f|\leqslant 1$. Therefore

$$f\circ\beta(t)\leqslant f(q)+t$$

for any $t\geqslant 0$.

Set $\ell(t)=|p-\beta(t)|$. Applying 13.22*a* and the first variation inequality (6.6), we get

$$\begin{aligned}\ell^2(t)'&\leqslant 2\cdot\left(f\circ\beta(t)-f(p)\right)\leqslant\\&\leqslant 2\cdot\left(f(q)+t-f(p)\right).\end{aligned}$$

Therefore

$$\ell^2(t)-\ell^2(0)\leqslant 2\cdot\left(f(q)-f(p)\right)\cdot t+t^2.$$

It proves the needed inequality in case $s=0$. Combining it with the first distance estimate (16.12), we get the result in case $s\leqslant t$. The case $s\geqslant t$ follows by switching the roles of s and t.. ☐

Proof of 16.21. Consider two Busemann functions, bus_+ and bus_-, asociated with half-lines $\gamma:[0,\infty)\to\mathcal{L}$ and $\gamma:(-\infty,0]\to\mathcal{L}$ respectively; that is,

$$\mathrm{bus}_\pm(x)=\lim_{t\to\infty}|\gamma(\pm t)-x|-t.$$

According to Exercise 8.24, both functions $\mathrm{bus}_\pm$ are concave.

Fix $x\in\mathcal{L}$. Note that since γ is a line, we have

$$\mathrm{bus}_+(x)+\mathrm{bus}_-(x)\geqslant 0.$$

On the other hand, by 8.23*b*, $f(t)=\mathrm{dist}_x^2\circ\gamma(t)$ is 2-concave. In particular, $f(t)\leqslant t^2+at+b$ for some constants $a,b\in\mathbb{R}$. Passing to the limit as $t\to\pm\infty$, we have

$$\mathrm{bus}_+(x)+\mathrm{bus}_-(x)\leqslant 0.$$

Hence

$$\mathrm{bus}_+(x)+\mathrm{bus}_-(x)=0$$

for any $x\in\mathcal{L}$. In particular the functions $\mathrm{bus}_\pm$ are affine; that is, they are convex and concave at the same time.

It follows that for any x,

$$\begin{aligned}|\nabla_x \operatorname{bus}_\pm| &= \sup\{\,\boldsymbol{d}_x \operatorname{bus}_\pm(\xi) : \xi \in \Sigma_x\,\} = \\ &= \sup\{\,-\boldsymbol{d}_x \operatorname{bus}_\mp(\xi) : \xi \in \Sigma_x\,\} \equiv \\ &\equiv 1.\end{aligned}$$

By Exercise 16.10, a 1-Lipschitz curve α such that $\operatorname{bus}_\pm(\alpha(t)) = t + c$ is a $\operatorname{bus}_\pm$-gradient curve. In particular, $\alpha(t)$ is a bus_+-gradient curve if and only if $\alpha(-t)$ is a bus_--gradient curve. It follows that for any $t > 0$, the $\operatorname{bus}_\pm$-gradient flows commute; that is,

$$\mathrm{Flow}^t_{\operatorname{bus}_+} \circ \mathrm{Flow}^t_{\operatorname{bus}_-} = \mathrm{id}_{\mathcal{L}}\,.$$

Setting

$$\mathrm{Flow}^t = \left[\begin{matrix}\mathrm{Flow}^t_{\operatorname{bus}_+} & \text{if } t \geqslant 0\\ \mathrm{Flow}^t_{\operatorname{bus}_-} & \text{if } t < 0\end{matrix}\right.$$

defines an $\mathbb{R}$-action on $\mathcal{L}$.

Consider the level set $\mathcal{L}' = \operatorname{bus}_+^{-1}(0) = \operatorname{bus}_-^{-1}(0)$; it is a closed convex subset of $\mathcal{L}$, and therefore forms an Alexandrov space. Consider the map $h\colon \mathcal{L}' \times \mathbb{R} \to \mathcal{L}$ defined by $h\colon (x,t) \mapsto \mathrm{Flow}^t(x)$. Note that h is onto. Applying Lemma 16.25 for $\mathrm{Flow}^t_{\operatorname{bus}_+}$ and $\mathrm{Flow}^t_{\operatorname{bus}_-}$ shows that h is short and non-contracting at the same time; that is, h is an isometry. □

E Radial curves

The radial curves are specially reparametrized gradient curves for distance functions. This parametrization makes them behave like unit-speed geodesics in a natural comparison sense (16F).

16.26. Definition. *Assume $\mathcal{L}$ is a complete length* CBB *space, $\kappa \in \mathbb{R}$, and $p \in \mathcal{L}$. A curve*

$$\sigma\colon [s_{\min}, s_{\max}) \to \mathcal{L}$$

is called a (p,κ)-radial curve if

$$s_{\min} = |p - \sigma(s_{\min})| \in (0, \tfrac{\varpi^\kappa}{2})$$

and σ satisfies the differential equation

❶ $$\sigma^+(s) = \frac{\operatorname{tg}^\kappa|p - \sigma(s)|}{\operatorname{tg}^\kappa s}\cdot\nabla_{\sigma(s)} \operatorname{dist}_p$$

for any $s \in [s_{\min}, s_{\max})$, where $\operatorname{tg}^\kappa x := \frac{\operatorname{sn}^\kappa x}{\operatorname{cs}^\kappa x}$.

If $x = \sigma(s_{\min})$, we say that σ starts at x.

Note that according to the definition, $s_{\max} \leqslant \frac{\varpi^\kappa}{2}$.

In the remainder of the chapter we will see that (p,κ)-radial curves work best for CBB(κ) spaces.

16.27. Definition. *Let $\mathcal{L}$ be a complete length* CBB *space and $p \in \mathcal{L}$. A unit-speed geodesic $\gamma\colon \mathbb{I} \to \mathcal{L}$ is called a p-radial geodesic if $|p - \gamma(s)| \equiv s$.*

The proofs of the following two propositions follow directly from the definitions.

16.28. Proposition. *Let $\mathcal{L}$ be a complete length* CBB *space and $p \in \mathcal{L}$. Assume $\frac{\varpi^\kappa}{2} \geqslant s_{\max}$. Then any p-radial geodesic $\gamma\colon [s_{\min}, s_{\max}) \to \mathcal{L}$ is a (p,κ)-radial curve.*

16.29. Proposition. *Suppose $\mathcal{L}$ is a complete length* CBB *space, $p \in \mathcal{L}$ and $\sigma\colon [s_{\min}, s_{\max}) \to \mathcal{L}$ is a (p,κ)-radial curve. Then for any $s \in [s_{\min}, s_{\max})$, we have $|p - \sigma(s)| \leqslant s$.*

Moreover, if for some s_0 we have $|p - \sigma(s_0)| = s_0$, then the restriction $\sigma|_{[s_{\min}, s_0]}$ is a p-radial geodesic.

16.30. Existence and uniqueness. *Let $\mathcal{L}$ be a complete length* CBB *space, $\kappa \in \mathbb{R}$, $p \in \mathcal{L}$, and $x \in \mathcal{L}$. Assume $0 < |p - x| < \frac{\varpi^\kappa}{2}$. Then there is a unique (p,κ)-radial curve $\sigma\colon [|p - x|, \frac{\varpi^\kappa}{2}) \to \mathcal{L}$ that starts at x; that is, $\sigma(|p - x|) = x$.*

Proof; existence. Let

$$\mathrm{itg}^\kappa\colon [0, \tfrac{\varpi^\kappa}{2}) \to \mathbb{R}, \quad \mathrm{itg}^\kappa(t) = \int_0^t \mathrm{tg}^\kappa \xi \cdot d\xi.$$

Clearly itg^κ is smooth and increasing. From 5.18 it follows that the composition

$$f = \mathrm{itg}^\kappa \circ \mathrm{dist}_p$$

is semiconcave in $\mathrm{B}(p, \frac{\varpi^\kappa}{2})$.

According to 16.15, there is an f-gradient curve $\alpha\colon [0, t_{\max}) \to \mathcal{L}$ defined on the maximal interval such that $\alpha(0) = x$.

Now consider a solution of the differential equation for $\tau(t)$, $\tau' = (\mathrm{tg}^\kappa \tau)^2$, $\tau(0) = r$. Note that $\tau(t)$ is also a gradient curve for the function itg^κ defined on $[0, \frac{\varpi^\kappa}{2})$. Direct calculations show that the composition $\alpha \circ \tau^{-1}$ is a (p,κ)-radial curve.

Uniqueness. Assume σ^1, σ^2 are two (p,κ)-radial curves that start at x. Then the compositions $\sigma^i \circ \tau$ both give f-gradient curves. By Picard's theorem (16.14), we have $\sigma^1 \circ \tau \equiv \sigma^2 \circ \tau$. Therefore $\sigma^1(s) = \sigma^2(s)$ for any $s \geqslant r$ such that both sides are defined. ☐

F Radial comparisons

In this section we show that radial curves behave in a comparison sense like unit-speed geodesics.

16.31. Radial monotonicity. *Let $\mathcal{L}$ be a complete length* $\mathrm{CBB}(\kappa)$ *space and p, q be distinct points in $\mathcal{L}$. Let $\sigma\colon [s_{\min}, \frac{\varpi^\kappa}{2}) \to \mathcal{L}$ be a (p,κ)-radial curve. Then the function*

$$s \mapsto \tilde{\measuredangle}^\kappa\{|q-\sigma(s)|; |p-q|, s\}$$

is nonincreasing in its entire domain of definition.

Radial monotonicity implies the following by straightforward calculations.

16.32. Corollary. *Let $\kappa \leqslant 0$, $\mathcal{L}$ be a complete* $\mathrm{CBB}(\kappa)$ *space, and $p, q \in \in \mathcal{L}$. Let $\sigma\colon [s_{\min}, \frac{\varpi^\kappa}{2}) \to \mathcal{L}$ be a (p,κ)-radial curve. Then for any $w \geqslant 1$, the function*

$$s \mapsto \tilde{\measuredangle}^\kappa\{|q-\sigma(s)|; |p-q|, w{\cdot}s\}$$

is nonincreasing in its entire domain of definition.

16.33. Radial comparison. *Let $\mathcal{L}$ be a complete length* $\mathrm{CBB}(\kappa)$ *space and $p \in \mathcal{L}$. Let $\rho\colon [r_{\min}, \frac{\varpi^\kappa}{2}) \to \mathcal{L}$ and $\sigma\colon [s_{\min}, \frac{\varpi^\kappa}{2}) \to \mathcal{L}$ be two (p,κ)-radial curves. Let*

$$\varphi_{\min} = \tilde{\measuredangle}^\kappa\left(p\,{}^{\rho(r_{\min})}_{\sigma(s_{\min})}\right).$$

Then for any $r \in [r_{\min}, \frac{\varpi^\kappa}{2})$ and $s \in [s_{\min}, \frac{\varpi^\kappa}{2})$, we have

$$\tilde{\measuredangle}^\kappa\{|\rho(r)-\sigma(s)|; r, s\} \leqslant \varphi_{\min},$$

or equivalently,

$$|\rho(r)-\sigma(s)| \leqslant \tilde{\curlyvee}^\kappa\{\varphi_{\min}; r, s\}.$$

We prove Theorems 16.31 and 16.33 simultaneously. The proof is an application of 13.22 plus trigonometric manipulations. We give a proof first in the simplest case $\kappa = 0$, since that is easier to follow, and then in the harder case $\kappa \neq 0$. The arguments for both cases are nearly the same, but the case $\kappa \neq 0$ requires an extra idea.

Proof of 16.31 and 16.33 in case $\kappa = 0$. Set

$$\begin{aligned} R &= R(r) = |p-\rho(r)|,\\ S &= S(s) = |p-\sigma(s)|,\\ \ell &= \ell(r,s) = |\rho(r)-\sigma(s)|,\\ \varphi &= \varphi(r,s) = \tilde{\measuredangle}^0\{\ell(r,s); r, s\}. \end{aligned}$$

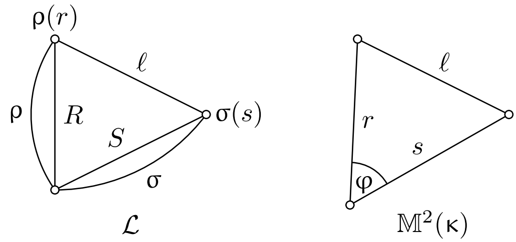

It will be sufficient to prove the following inequalities:

$$(*)^0_\varphi \qquad \tfrac{\partial^+}{\partial r}\varphi(s_{\min}, r) \leqslant 0, \qquad \tfrac{\partial^+}{\partial s}\varphi(s, r_{\min}) \leqslant 0$$

$$(**)^0_\varphi \qquad s\cdot\tfrac{\partial^+}{\partial s}\varphi + r\cdot\tfrac{\partial^+}{\partial r}\varphi \leqslant 0.$$

The radial monotonicity follows from $(*)^0_\varphi$. The radial comparison follows from $(*)^0_\varphi$ and $(**)^0_\varphi$. Indeed, one can connect $(s_{\min}, r_{\min})$ and (s_0, r_0) in $[s_{\min}, \infty) \times [r_{\min}, \infty)$ by a join of a coordinate segment and a segment defined by $r/s = r_0/s_0$ as in the figure. According to $(*)^0_\varphi$ and $(**)^0_\varphi$, φ does not increase while the pair (r, s) moves along this join with nondecreasing r and s. Thus $\varphi(r_0, s_0) \leqslant \varphi(r_{\min}, s_{\min}) = \varphi_{\min}$.

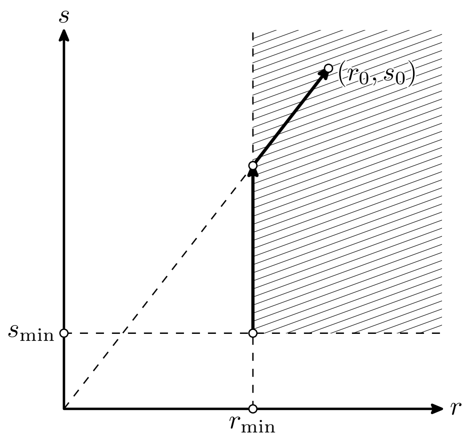

It remains to show $(*)^0_\varphi$ and $(**)^0_\varphi$. First let us rewrite the inequalities $(*)^0_\varphi$ and $(**)^0_\varphi$ in an equivalent form:

$$(*)^0_\ell \qquad \begin{aligned} &\tfrac{\partial^+}{\partial s}\ell(s, r_{\min}) \leqslant \cos\measuredangle^0\{r_{\min}; s, \ell\}, \\ &\tfrac{\partial^+}{\partial r}\ell(s_{\min}, r) \leqslant \cos\measuredangle^0\{s_{\min}; r, \ell\}, \end{aligned}$$

$$(**)^0_\ell \qquad s\cdot\tfrac{\partial^+}{\partial s}\ell + r\cdot\tfrac{\partial^+}{\partial r}\ell \leqslant s\cdot\cos\measuredangle^0\{r; s, \ell\} + r\cdot\cos\measuredangle^0\{s; r, \ell\} = \ell.$$

Let

$$(A)^0 \qquad f = \tfrac{1}{2}\cdot \mathrm{dist}_p^2 .$$

Clearly f is 1-concave, and

$$(B)^0 \qquad \rho^+(r) = \tfrac{1}{r}\cdot\nabla_{\rho(r)} f \quad\text{and}\quad \sigma^+(s) = \tfrac{1}{s}\cdot\nabla_{\sigma(s)} f.$$

Thus from 13.22, we have

$$(C)^0 \qquad \tfrac{\partial^+}{\partial r}\ell = -\tfrac{1}{r}\cdot\langle\nabla_{\rho(r)} f, \uparrow_{[\rho(r)\sigma(s)]}\rangle \leqslant \frac{\ell^2+R^2-S^2}{2\cdot\ell\cdot r}.$$

Since $R(r)\leqslant r$ and $S(s_{\min}) = s_{\min}$, then

$$(D)^0 \qquad \begin{aligned}\tfrac{\partial^+}{\partial r}\ell(r, s_{\min}) &\leqslant \frac{\ell^2+r^2-s_{\min}^2}{2\cdot\ell\cdot r} = \\ &= \cos\measuredangle^0\{s_{\min}; r, \ell\},\end{aligned}$$

which is the first inequality in $(*)^0_\ell$. By switching ρ and σ we obtain the second inequality in $(*)^0_\ell$. Further, adding $(C)^0$ and its mirror-inequality for $\frac{\partial^+}{\partial s}\ell$, we have

$$(E)^0 \qquad r\cdot\tfrac{\partial^+}{\partial r}\ell + s\cdot\tfrac{\partial^+}{\partial s}\ell \leqslant \frac{\ell^2+R^2-S^2}{2\cdot\ell} + \frac{\ell^2+S^2-R^2}{2\cdot\ell} = \ell,$$

namely $(**)^0_\ell$. □

Proof of 16.31 and 16.33 in case $\kappa\neq 0$. As before, let

$$\begin{aligned} R &= R(r) = |p-\rho(r)|, & \ell &= \ell(r,s) = |\rho(r)-\sigma(s)|,\\ S &= S(s) = |p-\sigma(s)|, & \varphi &= \varphi(r,s) = \measuredangle^\kappa\{\ell(r,s); r, s\}.\end{aligned}$$

It suffices to prove the following three inequalities:

$$(*)^\pm_\varphi \qquad \tfrac{\partial^+}{\partial r}\varphi(s_{\min}, r)\leqslant 0, \quad \tfrac{\partial^+}{\partial s}\varphi(s, r_{\min})\leqslant 0,$$

$$(**)^\pm_\varphi. \qquad \mathrm{sn}^\kappa s\cdot\mathrm{cs}^\kappa S\cdot\tfrac{\partial^+}{\partial s}\varphi + \mathrm{sn}^\kappa r\cdot\mathrm{cs}^\kappa R\cdot\tfrac{\partial^+}{\partial r}\varphi\leqslant 0.$$

Then radial monotonicity follows from $(*)^\pm_\varphi$. The radial comparison follows from $(*)^0_\varphi$ and $(**)^\pm_\varphi$. Indeed, the functions $s\mapsto \mathrm{sn}^\kappa s\cdot\mathrm{cs}^\kappa S$ and $r\mapsto\mathrm{sn}^\kappa r\cdot\mathrm{cs}^\kappa R$ are Lipschitz. Thus there is a solution for the differential equation

$$(r', s') = (\mathrm{sn}^\kappa s\cdot\mathrm{cs}^\kappa S, \mathrm{sn}^\kappa r\cdot\mathrm{cs}^\kappa R)$$

with any initial data $(r_0, s_0)\in[r_{\min}, \frac{\varpi^\kappa}{2})\times[s_{\min}, \frac{\varpi^\kappa}{2})$. (Unlike the case $\kappa = 0$, the solution cannot be written explicitly.) Since $\mathrm{sn}^\kappa s\cdot\mathrm{cs}^\kappa S$,

$\operatorname{sn}^\kappa r\cdot\operatorname{cs}^\kappa R>0$, this solution $(r(t),s(t))$ must meet one of the coordinate rays $\{r_{\min}\}\times[s_{\min},\frac{\varpi^\kappa}{2})$ or $[r_{\min},\frac{\varpi^\kappa}{2})\times\{s_{\min}\}$. That is, one can connect the pair $(s_{\min},r_{\min})$ to (s_0,r_0) by a join of a coordinate segment and part of the solution $(r(t),s(t))$. According to $(*)^\pm_\varphi$ and $(**)^\pm_\varphi$, the value of φ does not increase while the pair (r,s) moves along this join in direction of increasing r and s. Thus $\varphi(r_0,s_0)\leqslant\varphi(r_{\min},s_{\min})$.

As before, we rewrite the inequalities $(*)^\pm_\varphi$ and $(**)^\pm_\varphi$ in terms of ℓ:

$$(*)^\pm_\ell \qquad \begin{aligned}\tfrac{\partial^+}{\partial s}\ell(s,r_{\min})&\leqslant\cos\measuredangle^\kappa\{r_{\min};s,\ell\},\\ \tfrac{\partial^+}{\partial r}\ell(s_{\min},r)&\leqslant\cos\measuredangle^\kappa\{s_{\min};r,\ell\},\end{aligned}$$

$$(**)^\pm_\ell. \qquad \begin{aligned}&\operatorname{sn}^\kappa s\cdot\operatorname{cs}^\kappa S\cdot\tfrac{\partial^+}{\partial s}\ell+\operatorname{sn}^\kappa r\cdot\operatorname{cs}^\kappa R\cdot\tfrac{\partial^+}{\partial r}\ell\leqslant\\ &\qquad\leqslant\operatorname{sn}^\kappa s\cdot\operatorname{cs}^\kappa S\cdot\cos\measuredangle^\kappa\{r;s,\ell\}+\operatorname{sn}^\kappa r\cdot\operatorname{cs}^\kappa R\cdot\cos\measuredangle^\kappa\{s;r,\ell\}.\end{aligned}$$

Let

$$(A)^\pm \qquad f=-\tfrac{1}{\kappa}\cdot\operatorname{cs}^\kappa\circ\operatorname{dist}_p=\operatorname{md}^\kappa\circ\operatorname{dist}_p-\tfrac{1}{\kappa}.$$

Clearly $f''+\kappa\cdot f\leqslant0$ and

$$(B)^\pm \qquad \begin{aligned}\rho^+(r)&=\frac{1}{\operatorname{tg}^\kappa r\cdot\operatorname{cs}^\kappa R}\cdot\nabla_{\rho(r)}f,\\ \sigma^+(s)&=\frac{1}{\operatorname{tg}^\kappa s\cdot\operatorname{cs}^\kappa S}\cdot\nabla_{\sigma(s)}f.\end{aligned}$$

Thus from 13.22, we have

$$(C)^\pm \qquad \begin{aligned}\tfrac{\partial^+}{\partial r}\ell&=-\frac{1}{\operatorname{tg}^\kappa r\cdot\operatorname{cs}^\kappa R}\cdot\langle\nabla_{\rho(r)}f,\uparrow_{[\rho(r)\sigma(s)]}\rangle\leqslant\\ &\leqslant\frac{1}{\operatorname{tg}^\kappa r\cdot\operatorname{cs}^\kappa R}\cdot\frac{\operatorname{cs}^\kappa S-\operatorname{cs}^\kappa R\cdot\operatorname{cs}^\kappa\ell}{\kappa\cdot\operatorname{sn}^\kappa\ell}=\\ &=\frac{\frac{\operatorname{cs}^\kappa S}{\operatorname{cs}^\kappa R}-\operatorname{cs}^\kappa\ell}{\kappa\cdot\operatorname{tg}^\kappa r\cdot\operatorname{sn}^\kappa\ell}.\end{aligned}$$

Note that for all $\kappa\neq0$, the function $x\mapsto\frac{1}{\kappa\cdot\operatorname{cs}^\kappa x}$ is increasing. Thus, since $R(r)\leqslant r$ and $S(s_{\min})=s_{\min}$, we have

$$(D)^\pm \qquad \begin{aligned}\tfrac{\partial^+}{\partial r}\ell(r,s_{\min})&\leqslant\frac{\frac{\operatorname{cs}^\kappa s_{\min}}{\operatorname{cs}^\kappa r}-\operatorname{cs}^\kappa\ell}{\kappa\cdot\operatorname{tg}^\kappa r\cdot\operatorname{sn}^\kappa\ell}=\\ &=\frac{\operatorname{cs}^\kappa s_{\min}-\operatorname{cs}^\kappa\ell\cdot\operatorname{cs}^\kappa r}{\kappa\cdot\operatorname{sn}^\kappa r\cdot\operatorname{sn}^\kappa\ell}=\\ &=\cos\measuredangle^\kappa\{s_{\min};r,\ell\},\end{aligned}$$

which is the first inequality in $(*)^{\pm}_{\ell}$ for $\kappa\ne 0$. By switching ρ and σ we obtain the second inequality in $(*)^{\pm}_{\ell}$. Further, adding $(C)^{\pm}$ and its mirror-inequality for $\tfrac{\partial^+}{\partial s}\ell$, we have

$$(E)^{\pm}\qquad \begin{aligned}
&\operatorname{sn}^\kappa r\cdot\operatorname{cs}^\kappa R\cdot\tfrac{\partial^+}{\partial r}\ell+\operatorname{sn}^\kappa s\cdot\operatorname{cs}^\kappa S\cdot\tfrac{\partial^+}{\partial s}\ell\leqslant\\
&\qquad\leqslant\frac{\operatorname{cs}^\kappa S\cdot\operatorname{cs}^\kappa r-\operatorname{cs}^\kappa\ell\cdot\operatorname{cs}^\kappa R\cdot\operatorname{cs}^\kappa r}{\kappa\cdot\operatorname{sn}^\kappa\ell}+\\
&\qquad\quad+\frac{\operatorname{cs}^\kappa R\cdot\operatorname{cs}^\kappa s-\operatorname{cs}^\kappa\ell\cdot\operatorname{cs}^\kappa S\cdot\operatorname{cs}^\kappa s}{\kappa\cdot\operatorname{sn}^\kappa\ell}=\\
&\qquad=\operatorname{sn}^\kappa r\cdot\operatorname{cs}^\kappa R\cdot\frac{\operatorname{cs}^\kappa s-\operatorname{cs}^\kappa\ell\cdot\operatorname{cs}^\kappa r}{\kappa\cdot\operatorname{sn}^\kappa r\cdot\operatorname{sn}^\kappa\ell}+\\
&\qquad\quad+\operatorname{sn}^\kappa s\cdot\operatorname{cs}^\kappa S\cdot\frac{\operatorname{cs}^\kappa r-\operatorname{cs}^\kappa\ell\cdot\operatorname{cs}^\kappa s}{\kappa\cdot\operatorname{sn}^\kappa s\cdot\operatorname{sn}^\kappa\ell}=\\
&\qquad=\operatorname{sn}^\kappa r\cdot\operatorname{cs}^\kappa R\cdot\cos\tilde\measuredangle^\kappa\{r;s,\ell\}+\\
&\qquad\quad+\operatorname{sn}^\kappa s\cdot\operatorname{cs}^\kappa S\cdot\cos\tilde\measuredangle^\kappa\{s;r,\ell\},
\end{aligned}$$

which is $(**)^{\pm}_{\ell}$. ☐

16.34. Exercise. *Suppose $\mathcal{L}$ is a complete length* CBB(κ) *space and $x,y,z\in\mathcal{L}$. Assume $\tilde\measuredangle^\kappa(z\,{}^x_y)=\pi$. Show that there is a geodesic $[xy]$ that contains z. In particular, x can be connected to y by a minimizing geodesic. (Compare to Exercise 10.19.)*

G Gradient exponential map

Let $\mathcal{L}$ be a complete length CBB(κ) space, $p\in\mathcal{L}$, and $\xi\in\Sigma_p$. Consider a sequence of points $x_n\in\mathcal{L}$ such that $\uparrow_{[px_n]}\to\xi$. Let $r_n=|p-x_n|$, and let $\sigma_n\colon[r_n,\tfrac{\varpi^\kappa}{2})\to\mathcal{L}$ be the (p,κ)-radial curve that starts at x_n.

By the radial comparison (16.33), the curves $\sigma_n\colon[r_n,\tfrac{\varpi^\kappa}{2})\to\mathcal{L}$ converge to a curve $\sigma_\xi\colon(0,\tfrac{\varpi^\kappa}{2})\to\mathcal{L}$, and this limit is independent of the choice of the sequence x_n. Let $\sigma_\xi(0)=p$, and if $\kappa>0$ define

$$\sigma_\xi(\tfrac{\varpi^\kappa}{2})=\lim_{t\to\frac{\varpi^\kappa}{2}}\sigma_\xi(t).$$

The resulting curve σ_ξ will be called the (p,κ)-*radial curve in direction* ξ.

Let us define the *gradient exponential map* as

$$\operatorname{gexp}^\kappa_p\colon\overline{\mathrm{B}}[0,\tfrac{\varpi^\kappa}{2}]\subset\mathrm{T}_p\to\mathcal{L}\colon r\cdot\xi\mapsto\sigma_\xi(r).$$

Here are properties of radial curves reformulated in terms of the gradient exponential map:

16.35. Theorem. *Let $\mathcal{L}$ be a complete length* CBB(κ) *space. Then:*

a) If $p,q\in\mathcal{L}$ are points such that $|p-q|\leqslant\frac{\varpi^\kappa}{2}$, then for any geodesic $[pq]$ in $\mathcal{L}$ we have

$$\operatorname{gexp}_p^\kappa(\log[pq])=q.$$

b) For any $v,w\in\overline{\mathrm{B}}[0,\frac{\varpi^\kappa}{2}]\subset\mathrm{T}_p$,

$$|\operatorname{gexp}_p^\kappa v-\operatorname{gexp}_p^\kappa w|\leqslant\tilde\Upsilon^\kappa[0\,{}^v_w].$$

In other words, if we denote by $\mathcal{T}_p^\kappa$ the set $\overline{\mathrm{B}}[0,\frac{\varpi^\kappa}{2}]\subset\mathrm{T}_p$ equipped with the metric $|v-w|_{\mathcal{T}_p^\kappa}=\tilde\Upsilon^\kappa[0\,{}^v_w]$, then

$$\operatorname{gexp}_p^\kappa:\mathcal{T}_p^\kappa\to\mathcal{L}$$

is a short map.

c) Suppose $p,q\in\mathcal{L}$ and $|p-q|\leqslant\frac{\varpi^\kappa}{2}$. If $v\in\mathrm{T}_p$, $|v|\leqslant1$, and

$$\sigma(t)=\operatorname{gexp}_p^\kappa(t\cdot v),$$

then the function

$$s\mapsto\tilde\measuredangle^\kappa(\sigma|_0^s,q):=\tilde\measuredangle^\kappa\{|q-\sigma(s)|;|q-\sigma(0)|,s\}$$

is nonincreasing in its entire domain of definition.

Proof. Follows directly from the construction of $\operatorname{gexp}_p^\kappa$ and the radial comparison (16.33). □

Applying the theorem above together with 15.13*c*, we obtain the following.

16.36. Corollary. *Let $\mathcal{L}$ be an m-dimensional complete length* CBB(κ) *space, $p\in\mathcal{L}$, and $0<R\leqslant\frac{\varpi^\kappa}{2}$. Then there is a short map $f\colon\overline{\mathrm{B}}[R]_{\mathbb{M}^m(\kappa)}\to\mathcal{L}$ such that $\operatorname{Im}f=\overline{\mathrm{B}}[p,R]\subset\mathcal{L}$.*

16.37. Exercise. *Let $\mathcal{L}\subset\mathbb{E}^2$ be the Euclidean halfplane. Clearly $\mathcal{L}$ is a two-dimensional complete length* CBB(0) *space. Given a point $x\in\mathbb{E}^2$, denote by* $\operatorname{proj}(x)$ *the closest point to x on $\mathcal{L}$.*

Apply the radial comparison (16.33) to show that for any interior point $p\in\mathcal{L}$ and any $v\in\mathbb{R}^2$ we have

$$\operatorname{gexp}_p v=\operatorname{proj}(p+v).$$

16.38. Exercise. *Suppose x,p, and q are points in a complete length* CBB(κ) *space, and $x\in[pq[$. Show that there is a unique vector $v\in\mathrm{T}_p$ such that $\operatorname{gexp}_p v=x$.*

16.39. Exercise. *Let $\mathcal{L}$ be an m-dimensional complete length* CBB(κ) *space. Writing $\operatorname{rad}\mathcal{L}=R$, prove that there is a $(\operatorname{sn}^\kappa R)$-Lipschitz map $\Phi\colon\mathbb{S}^{m-1}\to\mathcal{L}$ such that $\operatorname{Im}\Phi\supset\partial\mathcal{L}$.*

H Remarks

Gradient flow on Riemannian manifolds

The gradient flow for general semiconcave functions on smooth Riemannian manifolds can be introduced with much less effort. To do this note that the distance estimates proved in the Section 16B can be proved in the same way for gradient curves of smooth semiconcave subfunctions. By the Greene–Wu lemma [55], given a λ-concave function f, a compact set $K \subset \operatorname{Dom} f$, and $\varepsilon > 0$ there is a smooth $(\lambda - \varepsilon)$-concave function that is ε-close to f on K. Hence one can apply smoothing and pass to the limit as $\varepsilon \to 0$. Note that by the second distance estimate (16.13), the limit curve obtained does not depend on the smoothing.

Gradient curves of a family of functions

Gradient flow can be extended to a family of functions. This type of flow was studied by Chanyoung Jun [74, 75], by Lucas Ferreira and Julio Valencia-Guevara [51], and by Alexander Mielke, Riccarda Rossi, and Giuseppe Savaré [94]. We will follow the simplified and generalized approach given by Alexander Lytchak and the third author [88], where an application related to this type of flow is given. The original motivation of Chanyoung Jun came from the study of pursuit-evasion problems. Another application of this type of flow comes from the fact that the optimal transport plan, or equivalently geodesics in the Wasserstein metric, can be described as gradient flow for a family of semiconcave functions. This observation was used by the third author to prove Alexandrov spaces with nonnegative curvature have nonnegative Ricci curvature in the sense of Lott–Villani–Sturm [106].

Suppose that $\mathcal{Z}$ is either CBB or CAT. Let f_t be a family of functions defined on open subsets $\operatorname{Dom} f_t$ of $\mathcal{Z}$. More precisely, we assume that the parameter t lies in a real interval $\mathbb{I}$ and

$$\Omega = \{\, (x,t) \in \mathcal{Z} \times \mathbb{I} : x \in \operatorname{Dom} f_t \,\}$$

is an open subset in $\mathcal{Z} \times \mathbb{I}$.

A family of functions f_t is called Lipschitz if the function $(x,t) \mapsto f_t(x)$ is L-Lipschitz for some constant L.

A family of functions f_t will be called semiconcave if the function $x \mapsto f_t(x)$ is λ-concave for each t. A family f_t is called locally semiconcave if for each $(p_0, t_0) \in \Omega$ there is a neighborhood Ω' and $\lambda \in \mathbb{R}$ such that the restriction of f_t to Ω' is semiconcave.

One cannot expect that a direct generalization of Definition 16.7 holds for every family of functions f_t; that is, gradient curves of a family f_t

cannot be defined as curves satisfying the equation $\alpha^+ = \nabla_\alpha f$.

For example, consider a 1-Lipschitz curve α in the real line. It is reasonable to assume that α is an f_t-gradient curve for the family $f_t(x) = -|x - \alpha(t)|$. (Indeed α can be realized as a limit of gradient curves for a family of functions obtained by smoothing f_t.) On the other hand, $\alpha^+(t)$ might be undefined, and even if it is defined, in general $\alpha^+(t) \neq 0$, while $\nabla_{\alpha(t)} f_t \equiv 0$.

Instead we define an f_t-gradient curve as a Lipschitz curve α that satisfies the following inequality for any point p, time t, and small $\varepsilon > 0$:

❶ $$\operatorname{dist}_p \circ \alpha(t+\varepsilon) \leqslant \operatorname{dist}_p \circ \alpha(t) - \varepsilon\cdot \boldsymbol{d}_{\alpha(t)} f_t(\uparrow_{[\alpha(t) p]}) + o(\varepsilon).$$

If there is no geodesic $[\alpha(t)\, p]$ then we impose no condition.

If $\alpha^+(t) = \nabla_{\alpha(t)} f_t$ for all t, then ❶ holds by the definition of gradient (13.17). On the other hand, the example above shows that the converse does not hold; that is, ❶ generalizes Definition 16.7. The defining inequality ❶ is closely related to the so-called evolution variational inequality [13, Theorem 4.0.4(iii)].

16.40. Distance estimate. *Let f_t and h_t be two families of λ-concave functions on a complete length space $\mathcal{Z}$, and $s \geqslant 0$. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT*. Assume f_t and h_t have common domain $\Omega \subset \mathcal{Z} \times \mathbb{R}$, and $|f_t(x) - h_t(x)| \leqslant s$ for any $(x,t) \in \Omega$. Assume $t \mapsto \alpha(t)$ and $t \mapsto \beta(t)$ are f_t- and h_t-gradient curves respectively defined on a common interval $t \in [a,b)$, and let $\ell(t) = |\alpha(t) - \beta(t)|_{\mathcal{Z}}$. If for all t, a minimizing geodesic $[\alpha(t)\,\beta(t)]$ lies in $\{\, x \in \mathcal{Z} : (x,t) \in \Omega \,\}$, then*

$$\ell'(t) \leqslant \lambda\cdot\ell(t) + 2\cdot s/\ell(t),$$

whenever the left-hand side is defined. Moreover,

$$\ell(t)^2 + \tfrac{2\cdot s}{\lambda} \leqslant (\ell(a)^2 + \tfrac{2\cdot s}{\lambda})\cdot e^{2\cdot\lambda\cdot(t-a)}.$$

In particular, these inequalities hold for any $t \in \mathbb{I}$ if $\Omega \supset B(p, 2\cdot r) \times \mathbb{I}$ and $\alpha(t), \beta(t) \in B(p,r)$ for any $t \in \mathbb{I}$.

Note that if $f_t = h_t$ then $s = 0$; in this case the second inequality can be written as

❷ $$\ell(t) \leqslant \ell(a)\cdot e^{\lambda\cdot(t-a)}.$$

In particular, this inequality implies uniqueness of the future of gradient curves with given initial data. This inequality also makes it possible to estimate the distance between two gradient curves for close functions. In particular, it implies convergence for f_t^n-gradient curves if a sequence of ℓ-Lipschitz and λ-concave families f_t^n converges uniformly as $n \to \infty$.

Proof of 16.40. Fix a time moment t and set $f = f_t$ and $h = h_t$. Let p be the midpoint of the geodesic $[\alpha(t)\beta(t)]$. Let $\gamma\colon [0,\ell(t)] \to \mathcal{Z}$ be an arclength parametrization of $[\alpha(t)\beta(t)]$. Note that $\boldsymbol{d}_{\alpha(t)}f(\uparrow_{[\alpha(t)p]})$ is the right derivative of $f\circ\gamma$ at 0 and $-\boldsymbol{d}_{\alpha(t)}h(\uparrow_{[\beta(t)p]})$ is the left derivative of $h\circ\gamma$ at $\ell(t)$. Since f and h are λ-concave,

$$\begin{aligned}
f\circ\beta(t) &\leqslant f\circ\alpha(t) + \ell(t)\cdot\boldsymbol{d}_{\alpha(t)}f(\uparrow_{[\alpha(t)p]}) + \tfrac{1}{2}\cdot\lambda\cdot\ell(t)^2,\\
h\circ\alpha(t) &\leqslant h\circ\beta(t) + \ell(t)\cdot\boldsymbol{d}_{\alpha(t)}h(\uparrow_{[\beta(t)p]}) + \tfrac{1}{2}\cdot\lambda\cdot\ell(t)^2.
\end{aligned}$$

Adding these inequalities and taking into account $|f(x) - h(x)| < s$ for any x, we conclude that

$$\boldsymbol{d}_{\alpha(t)}f(\uparrow_{[\alpha(t)p]}) + \boldsymbol{d}_{\alpha(t)}h(\uparrow_{[\beta(t)p]}) \geqslant \lambda\cdot\ell(t) + 2\cdot s/\ell(t).$$

Applying the triangle inequality and the definition of gradient curve at p, we obtain

$$\begin{aligned}
\ell(t+\varepsilon) &= |\alpha(t+\varepsilon) - \beta(t+\varepsilon)| \leqslant\\
&\leqslant |\alpha(t+\varepsilon) - p| + |\beta(t+\varepsilon) - p| \leqslant\\
&\leqslant |\alpha(t) - p| - \varepsilon\cdot\boldsymbol{d}_{\alpha(t)}f(\uparrow_{[\alpha(t)p]}) +\\
&\quad + |\beta(t+\varepsilon) - p| - \varepsilon\cdot\boldsymbol{d}_{\beta(t)}h(\uparrow_{[\beta(t)p]}) + o(\varepsilon) =\\
&= \ell(t) - \varepsilon\cdot(\lambda\cdot\ell(t) + 2\cdot s/\ell(t)) + o(\varepsilon)
\end{aligned}$$

for $\varepsilon > 0$. The first inequality follows.

Since α and β are Lipschitz, $t \mapsto \ell(t)$ is a Lipschitz function. By Rademacher's theorem, its derivative ℓ' is defined almost everywhere and satisfies the fundamental theorem of calculus. Therefore the first inequality implies the second. □

16.41. Proposition. *Suppose $\mathcal{Z}$ is a complete length space that is either* CBB *or* CAT. *Let f_t be a family of λ-concave functions for $t \in [a,b)$, where* $\operatorname{Dom} f_t \supset B(z, 2\cdot r)$ *for some fixed $z \in \mathcal{Z}$, $r > 0$ and any t.*

Let $\alpha\colon [a,b) \to B(z,r)$ be Lipschitz. Then α is an f_t-gradient curve if and only if

❸
$$\begin{aligned}
&\operatorname{dist}_p\circ\alpha(t+\varepsilon) \leqslant\\
&\quad\leqslant \operatorname{dist}_p\circ\alpha(t) - \varepsilon\cdot\left[\frac{f_t(p) - f_t\circ\alpha(t)}{|p - \alpha(t)|} - \tfrac{\lambda}{2}\cdot|p - \alpha(t)|\right] + o(\varepsilon)
\end{aligned}$$

for any $t \in [a,b)$ and $p \in B(z,r)\setminus\{\alpha(t)\}$.

Proof. Note that the geodesics $[\alpha(t)p]$ lie in $\operatorname{Dom} f_t$ for any t.

Since f_t is λ-concave, we have

$$\boldsymbol{d}_{\alpha(t)}f_t(\uparrow_{[\alpha(t)p]}) \geqslant \frac{f(p)-f\circ\alpha(t)}{|p-\alpha(t)|} - \tfrac{\lambda}{2}\cdot|p-\alpha(t)|.$$

Hence the only-if part follows.

Given $p \in \mathcal{Z}$ and t, consider a point $\bar p \in [\alpha(t)p]$. Applying ❸ for $\bar p$, and the triangle inequality, we have

$$\operatorname{dist}_p\circ\alpha(t+\varepsilon) \leqslant \operatorname{dist}_p\circ\alpha(t) - \varepsilon\cdot\left[\frac{f(\bar p)-f\circ\alpha(t)}{|\bar p-\alpha(t)|} - \tfrac{\lambda}{2}\cdot|\bar p-\alpha(t)|\right] + o(\varepsilon).$$

By taking $\bar p$ close to $\alpha(t)$, the value $\frac{f(\bar p)-f\circ\alpha(t)}{|\bar p-\alpha(t)|} - \frac{\lambda}{2}\cdot|\bar p-\alpha(t)|$ can be made arbitrarily close to $\boldsymbol{d}_{\alpha(t)}f_t(\uparrow_{[\alpha(t)p]})$. Therefore, given $\delta>0$, the inequality

$$\operatorname{dist}_p\circ\alpha(t+\varepsilon) \leqslant \operatorname{dist}_p\circ\alpha(t) - \varepsilon\cdot\boldsymbol{d}_{\alpha(t)}f_t(\uparrow_{[\alpha(t)p]}) + \varepsilon\cdot\delta$$

holds for all sufficiently small positive values ε. Therefore ❶ holds. □

Now we are ready to formulate and prove global existence of gradient curves for time-dependent families — an analog of 16.17.

16.42. Theorem. *Suppose $\mathcal{Z}$ is a complete length space that is either* CBB *or* CAT. *Let $\{f_t\}$ be a family of functions defined on an open set*

$$\Omega = \{\,(x,t)\in\mathcal{Z}\times\mathbb{R} : x\in\operatorname{Dom} f_t\,\}.$$

Suppose that f_t is Lipschitz and locally semiconcave. Then for any time a and initial point $p\in\operatorname{Dom} f_a$, there is a unique f_t-gradient curve $t\mapsto\alpha(t)$ defined on a maximal semiopen interval $[a,b)$. Moreover, if $b<\infty$ then $(\alpha(t),t)$ escapes from any closed set $K\subset\Omega$.

Proof. Let L be a Lipschitz constant of f_t. Fix $b>a$ sufficiently small that $\operatorname{Dom} f_t \supset B(p,\varepsilon\cdot L)$ for any $t\in[a,b)$. Consider a sequence $a=t_0<t_1\dots$ $\dots<t_n=b$, and a piecewise constant family of functions on $B(p,\varepsilon\cdot L)$ defined by $\hat f_t = f_{t_i}$ if $t_i\leqslant t<t_{i+1}$.

Note that $\hat f_t$ is time-independent on each interval $[t_i,t_{i+1})$. By 16.17 applied recursively on each interval $[t_i,t_{i+1})$, the proposition holds for $\hat f_t$. That is, there is a unique $\hat f_t$-gradient curve $\hat\alpha$ that starts at p and is defined on the interval $[a,b)$.

The distance estimates (16.40) show that as the partition gets finer, the gradient curves $\hat\alpha$ form a Cauchy sequence; denote its limit by α.

Then

$$\begin{aligned}\operatorname{dist}_p \circ \hat\alpha(t+\varepsilon) \leqslant{}& \operatorname{dist}_p \circ \hat\alpha(t) - \\ &- \varepsilon\cdot\left[\frac{\hat f_t(p) - \hat f_t\circ\hat\alpha(t)}{|p-\alpha(t)|} - \tfrac{\lambda}{2}\cdot|p-\hat\alpha(t)|\right] + o(\varepsilon) \leqslant \\ \leqslant{}& \operatorname{dist}_p \circ \hat\alpha(t) - \\ &- \varepsilon\cdot\left[\frac{f_t(p) - f_t\circ\hat\alpha(t) - 2\cdot s}{|p-\alpha(t)|} - \tfrac{\lambda}{2}\cdot|p-\hat\alpha(t)|\right] + o(\varepsilon),\end{aligned}$$

where

$$s = \sup_{t,x}\{|f_t(x) - \hat f_t(x)|\}.$$

Since $s \to 0$ as $\hat\alpha \to \alpha$, then ❸ holds for α; that is, α is an f_t-gradient curve.

This proves short time existence. Applying this argument recursively, we can find a gradient curve defined on a maximal interval $[a,b)$. Uniqueness of this curve follows from the distance estimate ❷.

Note that α is L-Lipschitz. In particular, if $b<\infty$ then $\alpha(t)\to p'$ as $t\to b$. If $(p',b)\in\Omega$ then we can repeat the procedure; otherwise α escapes from any closed set in Ω. □

Gradient curves for non-Lipschitz functions

In this book, we only consider gradient curves for locally Lipschitz semiconcave subfunctions; this turns out to be sufficient for all our needs. However, instead of Lipschitz semiconcave subfunctions, it is more natural to consider upper semicontinuous semiconcave functions with target in $[-\infty,\infty)$, and to assume in addition that the functions take finite values at a dense set in the domain of definition. Suppose that $\mathcal{Z}$ is a complete length space that is either CBB or CAT. The set of such subfunctions on $\mathcal{Z}$ will be denoted by $\mathrm{LCC}(\mathcal{Z})$ (for **l**ower semi-**c**ontinous and semi-**c**oncave).

In this section we describe the adjustments needed to construct gradient curves for the subfunctions in $\mathrm{LCC}(\mathcal{Z})$.

This type of function appears in entropy and some other closely related functionals on the Wasserstein space over a CBB(0) space. The gradient flow for these functions plays an important role in the theory of optimal transport, see [127] and references there in.

Differential. First we need to extend the definition of differential (6.13) to LCC subfunctions.

Let $\mathcal{Z}$ be a complete length space and $f\in\mathrm{LCC}(\mathcal{Z})$. Suppose that $\mathcal{Z}$ is either CBB or CAT. Given a point $p\in\operatorname{Dom} f$ and a geodesic direction

$\xi = \uparrow_{[pq]}$, let $\hat{\boldsymbol{d}}_p f(\xi) = (f \circ \mathrm{geod}_{[pq]})^+(0)$. Since f is semiconcave, the value $\hat{\boldsymbol{d}}_p f(\xi)$ is defined if $f \circ \mathrm{geod}_{[pq]}(t)$ is finite at all sufficiently small values $t > 0$, but $\hat{\boldsymbol{d}}_p f(\xi)$ may take value ∞. Note that $\hat{\boldsymbol{d}}_p f$ is defined on a dense subset of Σ_p.

Let

$$\boldsymbol{d}_p f(\zeta) = \varlimsup_{\xi \to \zeta} \hat{\boldsymbol{d}}_p f(\xi),$$

and $\boldsymbol{d}_p f(v) = |v| \cdot \boldsymbol{d}_p f(\xi)$ if $v = |v| \cdot \xi$ for some $\xi \in \Sigma_p$.

In other words, we define differential as the smallest upper semi-continuous positive-homogeneous function $\boldsymbol{d}_p f\colon \mathrm{T}_p \to \mathbb{R}$ such that if $\hat{\boldsymbol{d}}_p f(\xi)$ is defined, then $\boldsymbol{d}_p f(\xi) \geqslant \hat{\boldsymbol{d}}_p f(\xi)$.

Existence and uniqueness of the gradient. Note that in the proof of 13.19, we used the Lipschitz condition just once, to show that

$$\begin{aligned} s &= \sup \{ (\boldsymbol{d}_p f)(\xi) : \xi \in \Sigma_p \} = \\ &= \varlimsup_{x \to p} \frac{f(x) - f(p)}{|x - p|} < \\ &< \infty. \end{aligned}$$

The value s above will be denoted by $|\nabla|_p f$. Note that if the gradient $\nabla_p f$ is defined then $|\nabla|_p f = |\nabla_p f|$, and otherwise $|\nabla|_p f = \infty$.

Summarizing the discussion above, we have the following.

13.19′ Existence and uniqueness of the gradient. *Assume $\mathcal{Z}$ is a complete space and $f \in \mathrm{LCC}(\mathcal{Z})$. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT. *Then for any point $p \in \mathrm{Dom}\, f$, either there is a unique gradient $\nabla_p f \in \mathrm{T}_p$ or $|\nabla|_p f = \infty$.*

Further, in all the results of Section 13E we may assume only that both f and the gradient of f are defined at the points under consideration. The proofs are the same.

Sections 13E–16B require almost no changes. Mainly, where appropriate one needs to change $|\nabla_p f|$ to $|\nabla|_p f$ and/or assume that the gradient is defined at the points of interest. Also ❶ in Theorem 16.3 is taken as the definition of gradient-like curve. Then the theorem states that any gradient-like curve $\alpha\colon \mathbb{I} \to \mathcal{Z}$ satisfies Definition 16.2 at $t \in \mathbb{I}$ if $\nabla_{\hat{\alpha}(s)} f$ is defined. Further, Definition 16.7, should be changed to the following:

16.7′. Definition. *Let $\mathcal{Z}$ be a complete length space and $f \in \mathrm{LCC}(\mathcal{Z})$. Suppose that $\mathcal{Z}$ is either* CBB *or* CAT.

A curve $\alpha\colon [t_{\min}, t_{\max}) \to \mathrm{Dom}\, f$ will be called an f-gradient curve if

$$\alpha^+(t) = \nabla_{\alpha(t)} f$$

iwhen $\nabla_{\alpha(t)} f$ *is defined and*

$$(f \circ \alpha)^+(t) = \infty$$

otherwise.

In the proof of local existence (16.15), condition (ii) should be changed to the following condition:

(ii)′ $f \circ \hat{\alpha}_n(\bar{\sigma}_i) - f \circ \hat{\alpha}_n(\sigma_i) > (\bar{\sigma}_i - \sigma_i) \cdot \max\{n, |\nabla|_{\hat{\alpha}_n(\sigma_i)} f - \frac{1}{n})\}$.

Any gradient curve $\alpha[0, \ell) \to \mathcal{Z}$ for a subfunction $f \in \mathrm{LCC}(\mathcal{Z})$ satisfies the equation

$$\alpha^+(t) = \nabla_{\alpha(t)} f$$

at all values t, with the possible exception of $t = 0$. In particular, the gradient of f is defined at all points of any f-gradient curve, with the exception of the initial point.

Slower radial curves

Let $\kappa \geqslant 0$. Assume that for some function ψ, the curves defined by the equation

$$\sigma^+(s) = \psi(s, |p - \sigma(s)|) \cdot \nabla_{\sigma(s)} \operatorname{dist}_p$$

satisfy radial comparison 16.33. Then in fact the $\sigma(s)$ are radial curves; that is,

$$\psi(s, |p - \sigma(s)|) = \frac{\operatorname{tg}^\kappa |p - \sigma(s)|)}{\operatorname{tg}^\kappa s},$$

see exercise 16.37.

In case $\kappa < 0$, such a function ψ is not unique. In particular, one can take curves defined by the simpler equation

$$\sigma^+(s) = \frac{\operatorname{sn}^\kappa |p - \sigma(s)|}{\operatorname{sn}^\kappa s} \cdot \nabla_{\sigma(s)} \operatorname{dist}_p = \frac{1}{\operatorname{sn}^\kappa s} \cdot \nabla_{\sigma(s)} (\operatorname{md}^\kappa \circ \operatorname{dist}_p).$$

Among all curves of that type, the radial curves for curvature κ as defined in 16.26 maximize the growth of $|p - \sigma(s)|$.

Radial curves for sets

Here we generalize the constructions of radial curves and gradient exponent. We show that one can use a distance function dist_A to any closed set A instead of the distance function to one point. We only give the corresponding definitions and state the results. The proofs are straightforward generalizations of the corresponding one-point-set version.

First we give a more general form of the definitions of radial curves (16.26) and radial geodesics (16.27):

16.43. Definition. *Let $\mathcal{L}$ be a complete length* CBB *space, $\kappa \in \mathbb{R}$, and $A \subset \mathcal{L}$ be a closed subset of $\mathcal{L}$. A curve $\sigma\colon [s_{\min}, s_{\max}) \to \mathcal{L}$ is called an (A,κ)-radial curve if $s_{\min} = \operatorname{dist}_A \sigma(s_{\min}) \in (0, \frac{\varpi^\kappa}{2})$, and σ satisfies the differential equation*

$$\sigma^+(s) = \frac{\operatorname{tg}^\kappa |p - \sigma(s)|}{\operatorname{tg}^\kappa s} \cdot \nabla_{\sigma(s)} \operatorname{dist}_A$$

for any $s \in [s_{\min}, s_{\max})$, where $\operatorname{tg}^\kappa x = \frac{\operatorname{sn}^\kappa x}{\operatorname{cs}^\kappa x}$.

If $x = \sigma(s_{\min})$, we say that σ starts at x.

16.44. Definition. *Let $\mathcal{L}$ be a complete length* CBB *space and $A \subset \subset \mathcal{L}$ be a closed subset of $\mathcal{L}$. A unit-speed geodesic $\gamma\colon \mathbb{I} \to \mathcal{L}$ is called an A-radial geodesic if $\operatorname{dist}_A \gamma(s) \equiv s$.*

The following propositions are analogous to 16.28 and 16.29. Their proofs follow directly from the definitions.

16.45. Proposition. *Let $\mathcal{L}$ be a complete length* CBB *space, $A \subset \mathcal{L}$ be a closed subset of $\mathcal{L}$. Assume that $\frac{\varpi^\kappa}{2} \geqslant s_{\max}$. Then any dist_A-radial geodesic $\gamma\colon [s_{\min}, s_{\max}) \to \mathcal{L}$ is an (A,κ)-radial curve.*

16.46. Proposition. *Let $\mathcal{L}$ be a complete length* CBB *space, $A \subset \mathcal{L}$ be a closed subset of $\mathcal{L}$, and $\sigma\colon [s_{\min}, s_{\max}) \to \mathcal{L}$ be an (A,κ)-radial curve. Then for any $s \in [s_{\min}, s_{\max})$, we have $\operatorname{dist}_A \sigma(s) \leqslant s$.*

Moreover, if for some s_0 we have $\operatorname{dist}_A \sigma(s_0) = s_0$, then the restriction $\sigma|_{[s_{\min}, s_0]}$ is an A-radial geodesic.

Here is the corresponding generalization of existence and uniqueness for (A,κ)-radial curves; it can be proved in the same way as 16.30.

16.47. Existence and uniqueness. *Let $\mathcal{L}$ be a complete length* CBB *space, $\kappa \in \mathbb{R}$, $A \subset \mathcal{L}$ be a closed subset of $\mathcal{L}$, and $x \in \mathcal{L}$. Assume $0 < \operatorname{dist}_A x < \frac{\varpi^\kappa}{2}$. Then there is a unique (A,κ)-radial curve $\sigma\colon [\operatorname{dist}_A x, \frac{\varpi^\kappa}{2}) \to \mathcal{L}$ that starts at x.*

Next we formulate radial monotonicity and radial comparison for (A,κ)-radial curves. The proof of these two statements are almost exactly the same as the proofs of 16.31 and 16.33.

16.48. Radial monotonicity. *Let $\mathcal{L}$ be a complete length* CBB(κ) *space, $A \subset \mathcal{L}$ be a closed subset of $\mathcal{L}$, and $q \in \mathcal{L} \backslash A$. Assume $\sigma\colon [s_{\min}, \frac{\varpi^\kappa}{2}) \to \mathcal{L}$ is an (A,κ)-radial curve. Then the function*

$$s \mapsto \measuredangle^\kappa\{|q - \sigma(s)|; \operatorname{dist}_A q, s\}$$

is nonincreasing in its entire domain of definition.

To formulate a generalized radial comparison, we introduce a suitable notation. Given a subset A and two points x and y in a metric space, define

$$\measuredangle^\kappa \left(A\,{}^x_y\right) := \measuredangle^\kappa\{|x-y|; \operatorname{dist}_A x, \operatorname{dist}_A y\}.$$

Note that distances $|x-y|$, $\operatorname{dist}_A x$ and $\operatorname{dist}_A y$ might not satisfy the triangle inequality. Therefore the model angle $\measuredangle^\kappa\left(A\,{}^x_y\right)$ might be undefined even for $\kappa \leqslant 0$.

16.49. Radial comparison. *Let $\mathcal{L}$ be a complete length* CBB(κ) *space and $A \subset \mathcal{L}$ be a closed subset of $\mathcal{L}$. Assume $\rho\colon [r_{\min}, \frac{\varpi^\kappa}{2}) \to \mathcal{L}$ and $\sigma\colon [s_{\min}, \frac{\varpi^\kappa}{2}) \to \mathcal{L}$ are two dist_A-radial curves for curvature κ. Assume further that*

$$\varphi_{\min} = \measuredangle^\kappa\left(A\,{}^{\rho(r_{\min})}_{\sigma(s_{\min})}\right)$$

is defined. Then for any $r \in [r_{\min}, \frac{\varpi^\kappa}{2})$ and $s \in [s_{\min}, \frac{\varpi^\kappa}{2})$, we have

$$|\rho(r) - \sigma(s)| \leqslant \tilde{\Upsilon}^\kappa\{\varphi_{\min}; r, s\}.$$

Finally, suppose p is an isolated point of a closed subset A of a complete length CBB space $\mathcal{L}$. Applying the same limiting procedure as in Section 16G, for any $\xi \in \Sigma_p$ one can construct an (A,κ)-radial curve σ_ξ such that $\sigma_\xi(0) = p$ and $\sigma^+(0) = \xi$. Thus we obtain a map $\operatorname{gexp}^\kappa_A\colon \mathrm{T}_p \hookrightarrow \mathcal{L}$: $r\cdot\xi \mapsto \sigma_\xi(r)$. For this map, the following analog of 16.35 holds; the proof is straightforward.

16.50. Theorem. *Let $\mathcal{L}$ be a complete length* CBB(κ) *space, and $A \subset \mathcal{L}$ be a closed subset of A with an isolated point $p \in A$. Then:*

a) Let $\operatorname{dist}_A q = |p-q| \leqslant \frac{\varpi^\kappa}{2}$. Let $[pq]$ be an A-radial geodesic. Then

$$\operatorname{gexp}^\kappa_A(\log[pq]) = q.$$

b) For any $v, w \in \overline{\mathrm{B}}[0, \frac{\varpi^\kappa}{2}] \subset \mathrm{T}_p$,

$$|\operatorname{gexp}^\kappa_p v - \operatorname{gexp}^\kappa_p w| \leqslant \tilde{\Upsilon}^\kappa[0\,{}^v_w].$$

In other words, if we denote by $\mathcal{T}^\kappa_p$ the set $\overline{\mathrm{B}}[0, \frac{\varpi^\kappa}{2}] \subset \mathrm{T}_p$ equipped with the metric $|v-w|_{\mathcal{T}^\kappa_p} = \tilde{\Upsilon}^\kappa[0\,{}^v_w]$, then

$$\operatorname{gexp}^\kappa_p : \mathcal{T}^\kappa_p \to \mathcal{L}$$

is a short map.

c) Suppose $p, q \in \mathcal{L}$ and $|p-q| \leqslant \frac{\varpi^\kappa}{2}$. If $v \in \mathrm{T}_p$, $|v| \leqslant 1$ and

$$\sigma(t) = \operatorname{gexp}^\kappa_p(t\cdot v),$$

then the function $s \mapsto \measuredangle^\kappa(\sigma|^s_0, q)$ is nonincreasing in its entire domain of definition.

Appendix A

Semisolutions

1.3. Suppose α is a closed spherical curve. By Crofton's formula, the length of α is $\pi{\cdot}n_\alpha$, where n_α denotes the average number of crossings of α with equators.

Since α is closed, almost all equators cross it at an even number of points (we assume that ∞ is an even number). If $\operatorname{length}\alpha < 2{\cdot}\pi$ then $n_\alpha < 2$. Therefore there is an equator that does not cross α — hence the result.

2.9; *(a).* Note that any Cauchy sequence x_n in $(\mathcal{X}, \|{*}-{*}\|)$ is also Cauchy in $\mathcal{X}$. Since $\mathcal{X}$ is complete, x_n converges; denote its limit by x_∞.

Passing to a subsequence, we may assume that $\|x_{n-1} - x_n\| < \frac{1}{2^n}$. It follows that there is a 1-Lipschitz curve $\alpha\colon [0,1] \to (\mathcal{X}, \|{*}-{*}\|)$ such that $x_n = \alpha(\frac{1}{2^n})$ and $x_\infty = \alpha(0)$. In particular, $\|x_n - x_\infty\| \to 0$ and $n \to \infty$.

(b). Fix two points $x, y \in \mathcal{X}$ such that $\ell = \|x - y\| < \infty$. Let α_n be a sequence of paths from x to y such that $\operatorname{length}(\alpha_n) \to \ell$ as $n \to \infty$. Without loss of generality, we may assume that each α_n is $(\ell{+}1)$-Lipschitz.

Since $\mathcal{X}$ is compact, there is a partial limit α_∞ of α_n as $n \to \infty$. By semicontinuity of length, $\operatorname{length}\alpha_\infty \leqslant \ell$; that is; α is a shortest path in $\mathcal{X}$.

2.10. The following example was suggested by Fedor Nazarov [98].

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

$$\operatorname{length}_{\rho_1} \boldsymbol{x} = \int\limits_0^\ell \varphi \circ \boldsymbol{x}.$$

Note that the metric ρ_1 is bi-Lipschitz equivalent 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 endpoints, 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 endpoints. In particular, (B, ρ_1) has no geodesics.

2.15. The following example is taken from [25].

Consider the following subset of $\mathbb{R}^2$ equipped with the induced length metric

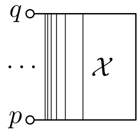

$$\mathcal{X} = \big((0,1] \times \{0,1\}\big) \cup \big(\{1, \tfrac{1}{2}, \tfrac{1}{3}, \dots\} \times [0,1]\big).$$

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 $\bar{\mathcal{X}}$. ☐

2.21 Let $\mathcal{X}$ be a compact metric space. Let us identify $\mathcal{X}$ with its image in $\operatorname{Bnd}(\mathcal{X}, \mathbb{R})$ under the Kuratowsky embedding (Section 2G). Denote by $\mathcal{K}$ the linear convex hull of $\mathcal{X}$ in the space of bounded functions on $\mathcal{X}$; that is, $x \in \mathcal{K}$ if and only if x cannot be separated from $\mathcal{X}$ by a hyperplane.

Since $\mathcal{X}$ is compact, so is $\mathcal{K}$. It remains to observe that $\mathcal{K}$ is a length space since it is convex.

3.10. Let $\mathcal{X}_n$ be the square $\{(x,y) \in \mathbb{R}^2, |x| \leqslant 1, |y| \leqslant 1\}$ with the metric induced by the ℓ^n-norm and let $f_n(x,y) = x$ for all n. Observe that $\mathcal{X}_\omega$ is the square with the metric induced by the ℓ^∞-norm where the limit function $f_\omega(x,y) = x$ is not concave.

4.13. Modify proof of 4.12, or apply 4.14*b*.

6.5. If $\measuredangle[p\,{}^{x}_{z}] + \measuredangle[p\,{}^{y}_{z}] < \pi$, then by the triangle inequality for angles (6.4) we have $\measuredangle[p\,{}^{x}_{y}] < \pi$. The latter implies that $[xy]$ fails to be minimizing near p.

6.8. By the definition of a right derivative, there is a geodesic γ such that both limits

$$\varlimsup_{\varepsilon \to 0+} \frac{|\alpha(\varepsilon) - \gamma(\varepsilon)|_{\mathcal{X}}}{\varepsilon} \quad \text{and} \quad \varlimsup_{\varepsilon \to 0+} \frac{|\beta(\varepsilon) - \gamma(\varepsilon)|_{\mathcal{X}}}{\varepsilon}$$

are arbitrarily small. Therefore

$$\varlimsup_{\varepsilon\to 0+}\frac{|\alpha(\varepsilon)-\beta(\varepsilon)|_{\mathcal{X}}}{\varepsilon}=0.$$

6.11. This follows directly from the definition.

6.12. Observe that

$$\operatorname{speed}_t\alpha=|\alpha^+(t)|=|\alpha^-(t)|.$$

Apply Theorem 5.10 to show that

$$|\alpha^+(t)-\alpha^-(t)|_{\mathrm{T}_{\alpha(t)}}=2\cdot\operatorname{speed}_t\alpha.$$

7.10. Choose two non-Euclidean norms $\|{*}\|_{\mathcal{X}}$ and $\|{*}\|_{\mathcal{Y}}$ on $\mathbb{R}^{10}$ such that the sum $\|{*}\|_{\mathcal{X}}+\|{*}\|_{\mathcal{Y}}$ is Euclidean. See [120] for more details.

8.3. Assume $|p-x^i|=|q-y^i|$ for each i. Observe and use that

$$|x^i-x^j|\leqslant|y^i-y^j|\quad\Longleftrightarrow\quad\tilde\measuredangle^\kappa(p\,{}^{x^i}_{x^j})\leqslant\tilde\measuredangle^\kappa(q\,{}^{y^i}_{y^j}).$$

8.4. Follows from the overlap lemma (10.2).

8.9. Modify the induced length metric on the unit sphere in an infinite-dimensional Hilbert space in small neighborhoods of a countable collection of points. To prove that the obtained space is CBB(0), you may need to use the technique from Halbeisen's example (13.6).

8.12. Mimic the proof of Theorem 8.11.

8.13. On the plane, any nonnegatively curved metric having an everywhere dense set of singular points will do the job, where by singular point we mean a point having total angle around it strictly smaller than $2\cdot\pi$.

Indeed, if x_i is a singular point, then there is $0<\varepsilon_i<1/20$ such that no geodesic with ends outside of $\mathrm{B}(x_i,r)$ can meet the ball $\mathrm{B}(x_i,\varepsilon_i\cdot r)$. The set

$$\Omega_n=\bigcup_i\mathrm{B}(x_i,\tfrac{\varepsilon_i}{n})$$

is open and everywhere dense. Note that Ω_n may intersect a geodesic of length $1/n$ only within $\frac{1}{10n}$ of its endpoints. The intersection of the Ω_n is a G-delta dense set that does not intersect the interior of any geodesic.

8.15. Note that rescaling does not change the space. Therefore if the space is CBB(κ) then it is CBB($\lambda\cdot\kappa$) for any $\lambda>0$. Passing to the limit as $\lambda\to 0$, we may assume that the space is CBB(0).

The point-on-side comparison (8.14b) for $p=v$, $x=w$, $y=-w$ and $z=0$ implies that

$$\|v+w\|^2+\|v-w\|^2\leqslant 2\cdot\|v\|^2+2\cdot\|w\|^2.$$

Applying the comparison for $p=v+w$, $x=w-v$, $y=v-w$ and $z=0$ gives the opposite inequality. That is, the parallelogram identity

$$\|v+w\|^2+\|v-w\|^2=2{\cdot}\|v\|^2+2{\cdot}\|w\|^2$$

holds for any vectors v and w. Whence the statement follows.

8.16. Apply the hinge comparison (8.14c).

8.18. Without loss of generality, we may assume that the points x, v, w, y appear on the geodesic $[xy]$ in that order. By the point-on-side comparison (8.14b) we have

$$\begin{aligned}\tilde\measuredangle^\kappa(x\,{}^y_p)\leqslant\tilde\measuredangle^\kappa(x\,{}^w_p)\leqslant\tilde\measuredangle^\kappa(x\,{}^v_p),\\ \tilde\measuredangle^\kappa(y\,{}^w_p)\geqslant\tilde\measuredangle^\kappa(y\,{}^v_p)\geqslant\tilde\measuredangle^\kappa(y\,{}^x_p).\end{aligned}$$

Therefore

$$\begin{aligned}\tilde\measuredangle^\kappa(x\,{}^y_p)<\tilde\measuredangle^\kappa(x\,{}^w_p)\quad&\Longrightarrow\quad\tilde\measuredangle^\kappa(x\,{}^y_p)<\tilde\measuredangle^\kappa(x\,{}^v_p),\\ \tilde\measuredangle^\kappa(y\,{}^x_p)<\tilde\measuredangle^\kappa(y\,{}^w_p)\quad&\Longleftarrow\quad\tilde\measuredangle^\kappa(y\,{}^x_p)<\tilde\measuredangle^\kappa(y\,{}^v_p).\end{aligned}$$

By Alexandrov's lemma (6.2), we have

$$\begin{aligned}\tilde\measuredangle^\kappa(x\,{}^y_p)<\tilde\measuredangle^\kappa(x\,{}^v_p)\quad&\Longleftrightarrow\quad\tilde\measuredangle^\kappa(y\,{}^x_p)<\tilde\measuredangle^\kappa(y\,{}^v_p),\\ \tilde\measuredangle^\kappa(x\,{}^y_p)<\tilde\measuredangle^\kappa(x\,{}^w_p)\quad&\Longleftrightarrow\quad\tilde\measuredangle^\kappa(y\,{}^x_p)<\tilde\measuredangle^\kappa(y\,{}^w_p).\end{aligned}$$

Hence the statement follows.

8.19. See the construction of Urysohn's space [56, $3.11\tfrac{3}{2}{}_+$].

8.20. Read [81].

8.21. Apply the angle-sidelength monotonicity (8.17) twice.

8.22. The first part follows from the angle-sidelength monotonicity (8.17). An example for the second part can be found among metrics on $\mathbb{R}^2$ induced by a norm. (Compare to Exercise 8.15.)

8.24 and 9.27. By the definition of Busemann function (see 6.1),

$$\begin{aligned}\exp(\sqrt{-\kappa}{\cdot}\,\mathrm{bus}_\gamma)&=\exp\left[\lim_{t\to\infty}\sqrt{-\kappa}{\cdot}(\mathrm{dist}_{\gamma(t)}-t)\right]=\\ &=\lim_{t\to\infty}\left(\exp\left[\sqrt{-\kappa}{\cdot}(\mathrm{dist}_{\gamma(t)}-t)\right]+\exp\left[\sqrt{-\kappa}{\cdot}(-\mathrm{dist}_{\gamma(t)}-t)\right]\right)=\\ &=\lim_{t\to\infty}\left(2{\cdot}\cosh\left[\sqrt{-\kappa}{\cdot}\,\mathrm{dist}_{\gamma(t)}\right]\cdot\exp\left[\sqrt{-\kappa}{\cdot}(-t)\right]\right).\end{aligned}$$

By the function comparison definitions of CAT(κ) space (9.25b) or CBB(κ) space (8.23b), for any $p\in\mathcal{U}$ the function $f=\cosh\sqrt{-\kappa}\circ\mathrm{dist}_p$ satisfies $f''+\kappa{\cdot}f\geqslant0$ (respectively $f''+\kappa{\cdot}f\leqslant0$). The result follows.

8.31. Read [105].

8.45. If $\operatorname{diam}(\mathcal{L}/G) > \frac{\pi}{2}$, then for some $x \in \mathcal{L}$ we have

$$\sup\{\,\operatorname{dist}_{G\cdot x}(y) : y \in \mathcal{L}\,\} > \tfrac{\pi}{2}.$$

Use comparison to show that there is a unique point y^* that lies at maximal distance from the orbit $G\cdot x$. Observe that y^* is a fixed point.

8.46. This exercise is based on the main idea in [68].

Assume there are 4 such points x_1, x_2, x_3, x_4. Since the space $\mathcal{L}$ is CBB(1) it is also CBB(0). By the angle comparison, the sum of the angles in any geodesic triangle in an CBB(0) space is $\geqslant \pi$. Therefore the average of the $\measuredangle[x_i\,{}^{x_j}_{x_k}]$ is larger than $\frac{\pi}{3}$. On the other hand, since each x_i has space of directions $\leqslant \frac{1}{2}\cdot\mathbb{S}^n$ and the perimeter of any triangle in $\frac{1}{2}\cdot\mathbb{S}^n$ is at most π, the average of $\measuredangle[x_i\,{}^{x_j}_{x_k}]$ is at most $\frac{\pi}{3}$ — a contradiction.

9.4. Suppose that

$$\tilde\measuredangle^{\kappa}(x^0\,{}^{x^1}_{x^2}) + \tilde\measuredangle^{\kappa}(x^0\,{}^{x^2}_{x^3}) < \tilde\measuredangle^{\kappa}(x^0\,{}^{x^1}_{x^3}).$$

Show that

$$\tilde\measuredangle^{\kappa}(x^2\,{}^{x^0}_{x^1}) + \tilde\measuredangle^{\kappa}(x^2\,{}^{x^1}_{x^3}) + \tilde\measuredangle^{\kappa}(x^2\,{}^{x^3}_{x^0}) > 2\cdot\pi.$$

Conclude that one can take $p = x^2$.

9.5. This is analogous to Exercise 8.3.

9.6. Read [119]. (The original proof [19] is much longer and is harder to follow.)

An example for the second part of the problem can be found among 4-point metric spaces. It is sufficient to take a generic convex quadrangle and increase one of its diagonals slightly; it will still satisfy the inequality for all relabeling but will fail to meet 9.2.

9.9. Suppose that a geodesic $[px]$ is not extendable beyond x. We may assume that $|p-x| < \varpi^{\kappa}$; otherwise move p along the geodesic toward x.

By the uniqueness of geodesics (9.8), any point y in a neighborhood $\Omega \ni x$ is connected to p by a unique geodesic path; denote it by γ_y. Moreover, $h_t(y) = \gamma_y(t)$ defines a homotopy, called the geodesic homotopy, between the identity map of Ω and the constant map with value p.

Since $[px]$ is not extendable, $x \notin h_t(\Omega)$ for any $t < 1$. In particular, the local homology groups vanish at x — a contradiction.

9.10. Choose a sequence of directions ξ_n at p of both-sided local geodesics; denote by γ_n the corresponding geodesics. Since the space $\mathcal{U}$ is locally compact, we may pass to a converging subsequence of (γ_n); it's limit is a local geodesic by Corollary 9.22. Denote the limit by γ_∞ and its direction by ξ_∞. By comparison, ξ_∞ is a limit of (ξ_n).

9.16. Follow the solution in the 8.15, reversing all the inequalities.

9.17. It is sufficient to show that if v and y are midpoints of geodesics $[uw]$ and $[xz]$ in $\mathcal{U}$, then

$$|v - y| \leqslant \tfrac{1}{2}\cdot(|u - x| + |w - z|).$$

Denote by p the midpoint of $[uz]$. Applying the angle-sidelength monotonicity (9.15) twice, we have

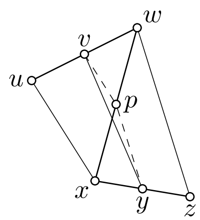

$$|v - p| \leqslant \tfrac{1}{2}\cdot|w - z|.$$

Similarly we have

$$|y - p| \leqslant \tfrac{1}{2}\cdot|u - x|.$$

It remains to add these two inequalities and apply the triangle inequality.

Comment. This inequality also follows directly from the majorization theorem (9.54).

9.20. The only-if part is evident.

Use 9.14 to show that $(c)\Rightarrow(b)\Rightarrow(a)$.

By 9.14, condition (a) implies that the natural map is distance-preserving on the sides $[\tilde{x}\tilde{y}]$ and $[\tilde{x}\tilde{z}]$. Applying it again, we have that condition (a) holds for all permutations of the labels x, y, z. Whence the natural map is distance-preserving on all three sides.

Comment. These conditions imply that the natural map can be extended to a distance-preserving map to the solid model triangle. In fact the image of the line-of-sight map (9.31) is isometric to the model triangle.

9.27. See the solution of Exercise 8.24.

9.72. Let $p \mapsto \bar{p}$ denote the closest-point projection to K. We need to show that $|\bar{p} - \bar{q}| \leqslant |p - q|$ for any $p, q \in \mathcal{U}$.

Assume $p \neq \bar{p} \neq \bar{q} \neq q$. Note that in this case $\measuredangle[\bar{p}\,{}^{p}_{\bar{q}}] \geqslant \tfrac{\pi}{2}$ and $\measuredangle[\bar{q}\,{}^{q}_{\bar{p}}] \geqslant \tfrac{\pi}{2}$. Otherwise a point on the geodesic $[\bar{p}\bar{q}]$ would be closer to p or to q than $\bar{p}$ or $\bar{q}$ respectively. The latter is impossible since K is convex and therefore $[\bar{p}\bar{q}] \subset K$.

Applying the arm lemma (9.59), we get the statement.

The cases $p = \bar{p} \neq \bar{q} \neq q$ and $p \neq \bar{p} \neq \bar{q} = q$ can be done similarly. The rest of the cases are trivial.

9.73. A more transparent, but less elementary solution via gradient flow is given by Alexander Lytchak and the third author [87].

Without loss of generality, we may assume that $p \in K$.

If $|K - x| \geqslant \pi$, then set $\Psi(x) = p$.

Otherwise, if $|K-x|<\pi$, by the closest-point projection lemma (9.71), there is a unique point $x^*\in K$ that minimizes distance to x; that is, $|x^*-x|=|K-x|$. Let us define ℓ_x, φ_x and ψ_x using the following identities:

$$\begin{aligned}\ell_x&=|p-x^*|,\\ \varphi_x&=\tfrac{\pi}{2}-|x^*-x|,\\ \sin\psi_x&=\sin\varphi_x\cdot\sin\ell_x,\quad 0\leqslant\psi_x\leqslant\tfrac{\pi}{2}.\end{aligned}$$

Let

$$\Psi(x)=\mathrm{geod}_{[px^*]}(\psi_x).$$

Note that Ψ is a retraction to K; that is, $\Psi(x)\in K$ for any $x\in\mathcal{U}$ and $\Psi(a)=a$ for any $a\in K$.

Let us show that Ψ is short. Given $x,y\in\mathrm{B}(K,\frac{\pi}{2})$, let

$$\begin{aligned}x'&=\Psi(x) & y'&=\Psi(y)\\ r&=|x-y| & r'&=|x'-y'|\\ d&=|x^*-y^*| & \alpha&=\tilde\measuredangle^1(p\,{}^{x^*}_{y^*}).\end{aligned}$$

Note that

❶ $$\cos r\leqslant\cos\varphi_x\cdot\cos\varphi_y-\cos d\cdot\sin\varphi_x\cdot\sin\varphi_y.$$

Indeed, if $x,y\notin K$, then $\measuredangle[x^*\,{}^{x}_{y*}],\measuredangle[y^*\,{}^{y}_{x*}]\geqslant\frac{\pi}{2}$ and the inequality ❶ follows from the arm lemma (9.59). If $x\in K$ and $y\notin K$, we obtain ❶ by the angle comparison (9.14*c*) since $\measuredangle[y^*\,{}^{y}_{x*}]\geqslant\frac{\pi}{2}$. In the same way, ❶ is proved if $x\notin K$ and $y\in K$. Finally, if $x,y\in K$, then $\varphi_x=\varphi_y=\frac{\pi}{2}$ and $r=d$; that is, the inequality trivially holds.

Further note that

$$\cos\alpha=\frac{\cos d-\cos\ell_x\cdot\cos\ell_y}{\sin\ell_x\cdot\sin\ell_y}.$$

Applying the angle-sidelength monotonicity (9.15), we have

$$\begin{aligned}\cos r'&\geqslant\cos\psi_x\cdot\cos\psi_y-\cos\alpha\cdot\sin\psi_x\cdot\sin\psi_y=\\ &=\cos\psi_x\cdot\cos\psi_y-(\cos d-\cos\ell_x\cdot\cos\ell_y)\cdot\sin\varphi_x\cdot\sin\varphi_y\geqslant\\ &\geqslant\cos\psi_x\cdot\cos\psi_y-\cos d\cdot\sin\varphi_x\cdot\sin\varphi_y.\end{aligned}$$

Note that $\psi_x\leqslant\varphi_x$ and $\psi_y\leqslant\varphi_y$; in particular,

$$\cos\varphi_x\cdot\cos\varphi_y\leqslant\cos\psi_x\cdot\cos\psi_y.$$

Hence

$$\cos r' \geqslant \cos r;$$

that is, the restriction $\Psi|_{\mathrm{B}(K,\frac{\pi}{2})}$ is short. Clearly Ψ is continuous. Since the complement of $\mathrm{B}(K,\frac{\pi}{2})$ is mapped to p, Ψ is short; that is,

❷ $$r' \leqslant r$$

for any $x, y \in \mathcal{U}$.

If we have equality in ❷ then

$$\cos \ell_x \cdot \cos \ell_y \cdot \sin \varphi_x \cdot \sin \varphi_y = 0.$$

If $K \subset \mathrm{B}(p,\frac{\pi}{2})$, then $\ell_x, \ell_y < \frac{\pi}{2}$, which implies that $x \in K$ or $y \in K$. Without loss of generality, we may assume that $x \in K$.

It remains to show that if $y \notin K$ then the inequality ❷ is strict. If $|K - y| \geqslant \frac{\pi}{2}$, then ❷ holds since the left-hand side is $< \frac{\pi}{2}$ while the right-hand side is $\geqslant \frac{\pi}{2}$. If $|K - y| < \frac{\pi}{2}$, then $\varphi_y > 0$. Clearly $\psi_y < \varphi_y$, hence the inequality ❷ is strict. □

Below you will find a geometric way to think about the given construction; in fact it is very close to the construction in the proof of Kirszbraun's theorem (10.14).

Geometric interpretation of the map Ψ. Let $\mathring{\mathcal{U}} = \mathrm{Cone}\,\mathcal{U}$, and denote by $\mathring{K}$ the subcone of $\mathring{\mathcal{U}}$ spanned by K. The space $\mathcal{U}$ can be naturally identified with the unit sphere in $\mathring{\mathcal{U}}$, that is, the set

$$\left\{\, z \in \mathring{\mathcal{U}} : |z| = 1 \,\right\}.$$

According to 11.7, $\mathring{\mathcal{U}}$ is CAT(0). Note that $\mathring{K}$ forms a convex closed subset of $\mathring{\mathcal{U}}$. According to 9.71, for any point x there is a unique point $\hat{x} \in \mathring{K}$ that minimizes the distance to x, that is, $|\hat{x} - x| = |K - x|$. (If $|\hat{x}| \neq 0$, then in the notation above we have $x^* = \frac{1}{|\hat{x}|} \cdot \hat{x}$.)

Consider the half-line $t \mapsto t \cdot p$ in $\mathring{\mathcal{U}}$. By comparison, for given $s \in \mathring{\mathcal{U}}$ the geodesics $\mathrm{geod}_{[s\ t\cdot p]}$ converge as $t \to \infty$ to a half-line, say $\alpha_s \colon [0,\infty) \to \mathring{\mathcal{U}}$.

Note that if $|x| = 1$, then $|\hat{x}| \leqslant 1$. By assumption, for any $a \in K$ the function $t \mapsto |\alpha_a(t)|$ is monotonically increasing. Therefore there is a unique value $t_x \geqslant 0$ such that $|\alpha_{\hat{x}}(t_x)| = 1$. Define $\Psi \colon \mathcal{U} \to K$ by

$$\Psi(x) = \alpha_{\hat{x}}(t_x).$$

9.32; *a*. Suppose that $x_n \to x_\infty$, $y_n \to y_\infty$ as $n \to \infty$, but $[x_n y_n]$ does not converge to $[x_\infty y_\infty]$. Since the space is proper, we can pass to a subsequence such that $[x_n y_n]$ converges to another geodesic. That is, we have at least two geodesics between x_∞ and y_∞.

b. The following example is taken from [25, Chapter I, Exercise 3.14].

Let Δ_n be a sequence of solid spherical triangles with angle $\frac{\pi}{4}$ and adjacent sides $\pi - \frac{1}{n}$. Let us glue each Δ_n to $[0,\pi]$ along an isometry of one of the longer sides. It remains to show that the obtained space $\mathcal{X}$ is a needed example.

9.40. Note that Q is CAT(0). Therefore by the Reshetnyak gluing theorem (9.38), by gluing γ_1 to γ_1' we obtain a CAT(0) space, say $\mathcal{U}'$.

The above curve resulting from gluing γ_1 to γ_1' forms together with γ_2 a both-sided infinite geodesic, say γ, in $\mathcal{U}'$. In particular, γ is a convex set isometric to $\mathbb{R}$.

Now glue Q along its boundary to $\mathcal{U}$ along γ. By the Reshetnyak gluing theorem, the resulting space is CAT(0).

It remains to note that this space can be obtained directly by gluing $\mathcal{U}$ to Q along corresponding half-lines.

9.41. Suppose that A is not convex. Then there is a geodesic $[xy]$ with ends in A that does not lie in A completely. Note that $[xy]$ can be lifted to two different geodesics with the same ends in the doubling, and apply uniqueness of geodesics (9.13).

9.42. Since K is π-convex, it is CAT(1). By 11.7, the spherical suspension $\operatorname{Susp} K$ is CAT(1) as well. Let us glue $\operatorname{Susp} K$ to $\mathcal{U}$ along K; according to the Reshetnyak gluing theorem, the resulting space, say $\mathcal{U}'$, is CAT(1).

Consider the geodesic path $\gamma\colon [0,1]$ from p to a pole of the suspension in $\mathcal{U}'$. Set $K_t = \mathcal{U} \cap \overline{\mathrm{B}}[\gamma(t), \frac{\pi}{2}]$. By 9.26, K_t is π-convex for any t, and monotonicity of the family should be evident.

Remark. Note that by applying Sharafutdinov retraction to the family of convex sets provided by the exercise, we get a short strong deformation retraction from $\overline{\mathrm{B}}[p, \frac{\pi}{2}]$ to K; that is, there is a family of maps $\varphi_t\colon \overline{\mathrm{B}}[p, \frac{\pi}{2}] \to \overline{\mathrm{B}}[p, \frac{\pi}{2}]$ such that the function $t \mapsto |\varphi_t(x) - \varphi_t(y)|$ is non-increasing for any pair of points $x, y \in \overline{\mathrm{B}}[p, \frac{\pi}{2}]$, $\varphi_t(x) = x$ for any $x \in K$ and $\varphi_1(\overline{\mathrm{B}}[p, \frac{\pi}{2}]) = K$. Moreover we can assume that there is a family of short maps $\varphi_t\colon \overline{\mathrm{B}}[p, \frac{\pi}{2}]) \to K_t$ such that $\varphi_t(x) = x$ for any t and $x \in K_t$. This leads to another solution of Exercise 9.73.

9.57. *(Easier way.)* Let $(t,s) \mapsto \gamma_t(s)$ be the line-of-sight map for α with respect to $\alpha(0)$, and $(t,s) \mapsto \tilde\gamma_t(s)$ be the line-of-sight map for $\tilde\alpha$ with respect to $\tilde\alpha(0)$. Consider the map $F\colon \operatorname{Conv}\tilde\alpha \to \mathcal{U}$ such that $F\colon \tilde\gamma_t(s) \mapsto \gamma_t(s)$.

Show that F majorizes α and conclude that F is distance-preserving.

(Harder way.) Prove and apply the following statement together with the Majorization theorem.

◇ Let α and β be two convex curves in $\mathbb{M}^2(\kappa)$. Assume

$$\operatorname{length}\alpha = \operatorname{length}\beta < 2\cdot\varpi^\kappa$$

and there is a short bijecction $f\colon \alpha \to \beta$. Then f is an isometry.

9.58. Suppose that points p, x, q, y appear on the curve in that cyclic order. Assume that the geodesics $[pq]$ and $[xy]$ do not intersect. Use the argument in the proof of the majorization theorem (9.54) to show that in this case there are nonequivalent majorization maps.

Now we can assume that pairs of geodesics $[pq]$ and $[xy]$ intersect for all choices of points p, x, q, y on the curve in that cyclic order. Show that in this case the convex hull K of the curve is isometric to a convex figure.

Note that the composition of a majorization map and closest point projection to K is a majorization. Show and use that the boundary of a convex figure in the plane admits a unique majorization up to equivalence.

Comment. A typical rectifiable closed curve in a CAT(0) space can be majorized by more than one convex figure. There are two exceptions: (1) if the majorization map is distance-preserving, and (2) if the curve is formed by three sides of a geodesic triangle. Richard Bishop asked if there are no other exceptions.

9.60. Look at the four triples in x^1, x^2, x^3, x^4. By hypothesis, for each triple there is a triangle $\triangle^i$ in $\mathbb{E}^2$ whose sidelengths are the three distances between pairs. By the arm Lemma (9.59), for each triple, $\mathcal{U}$ contains an isometric copy of the corresponding solid triangle. Consider adjacent pairs of these solid triangles in $\mathcal{U}$, with common sides moving around the sides of the quadrangle in order. The fourth solid triangle has a common side with the first. Therefore the union of the four solid triangles contains a triangle $\triangle$ with vertex x^1 and its other two vertexes in the intersections of the solid triangle opposite x^1 with the two solid triangles with vertex x^1 respectively, and where $\triangle$ has the same sidelengths as a Euclidean triangle. Therefore $\mathcal{U}$ contains an isometric copy of the solid Euclidean triangle with those side lengths. Moving around the adjacent pairs of solid triangles in $\mathcal{U}$ corresponding to the $\triangle^i$ shows that $\mathcal{U}$ contains an isometric copy of the solid Euclidean quadrangle, as required.

9.66. If ℓ and m do not intersect, then the double cover $\mathcal{X}$ is not simply connected. In particular, by the Hadamard–Cartan theorem, $\mathcal{X}$ is not CAT(0).

If ℓ and m intersect then $\mathcal{X}$ is a cone over a double cover Σ of $\mathbb{S}^2$ branching at two pairs (x, y) and (v, w) of antipodal points. Suppose $|x - v|_{\mathbb{S}^2} = \ell < \frac{\pi}{2}$. Note that the inverse image of $[xv]_{\mathbb{S}^2}$ is a closed geodesic of length $4{\cdot}\ell < 2{\cdot}\pi$. Therefore, by the generalized Hadamard–Cartan theorem, Σ is not CAT(1). Hence $\mathcal{X}$ is not CAT(0) by Theorem 11.7 on curvature of cones.

9.67. Let us do the second part first. Assume A has nonempty interior. Note that the space $\tilde{\mathcal{U}}$ is simply connected and locally isometric to the doubling $\mathcal{W}$ of $\mathcal{U}$ in A; that is, any point in $\tilde{\mathcal{U}}$ has a neighborhood that is isometric to a neighborhood of a point in $\mathcal{W}$.

By the Reshetnyak gluing theorem (9.38), $\mathcal{W}$ is CAT(0). By the Hadamard–Cartan theorem (9.61), $\tilde{\mathcal{U}}$ is CAT(0).

For the general case, apply the above argument to a closed ε-neighborhood of A and pass to a limit as $\varepsilon \to 0$.

The first part of the problem follows since a geodesic is a convex set.

9.68. Prove that the angle comparison (9.14*c*) holds.

9.69. Mimic the proof of the Hadamard–Cartan theorem.

10.5. Note that it is sufficient to show that any finite set of points $x^1, \dots$ $\dots, x^n \in \mathcal{X}$ lies in an isometric copy of a Euclidean polyhedron.

Observe that $\mathcal{X}$ is CBB(0) and CAT(0) at the same time. Show that there is a unique point p that minimizes the sum $|p - x^1| + \dots + |p - - x^n|$. Note that the vectors $v^i = \mathrm{lg}_{[px^i]}$ lie in a linear subspace of T_p, where lg_{x_0} is the gradient logarithm map $\mathrm{lg}_{x_0} \colon \mathcal{L} \to \mathrm{T}_{x_0}$ corresponding to the gradient exponential gexp_{x_0} (see Section 16G). Moreover if K is the convex hull of v_i, then the origin of T_p lies in the interior of K relative to its affine hull. Finally observe that the exponential map is defined on all of K and is distance-preserving. The statement follows since the exponential map sends $v^i \mapsto x^i$ for each i.

10.6. The answers are $s \leqslant \sqrt{3}$ and $s \leqslant 2$ respectively.

The upper bound $s \leqslant \sqrt{3}$ follows from (2+2)-point comparison.

The Euclidean space works as an example if s is smaller than the large diagonal of the double pyramid with unit side (that is, if $s \leqslant 2\cdot\sqrt{2/3}$). Otherwise take the product $\mathcal{K} \times \mathbb{R}$ with a 2-dimensional cone for the CAT(0) case. For the CBB(0) case, the needed space can be constructed by doubling a proper convex set $K \subset \mathbb{E}^3$ in its boundary. We assume that the points correspond to vertices of a regular tetrahedron with 3 vertices on the boundary of K and one in its interior; this point corresponds to a pair of points in the doubling at distance s from each other.

Remarks. Tetsu Toyoda [126] showed that any 5-point metric space that satisfies the (2+2)-point comparison admits a distance-preserving map into a CAT(0) space. For spaces with more than 5 points the condition is unknown; for 6-point metric spaces, a conjecture is formulated in [83].

10.7. Choose a quadruple of points p, q, r, s. Suppose that it admits a distance-preserving embedding into some $\mathbb{M}^2(\mathrm{K})$ for some $\mathrm{K} \geqslant \kappa$. Then

$$\tilde\measuredangle^{\mathrm{K}}(p\,{}^{q}_{r}) + \tilde\measuredangle^{\mathrm{K}}(p\,{}^{r}_{s}) + \tilde\measuredangle^{\mathrm{K}}(p\,{}^{a}_{q}) \leqslant 2\cdot\pi.$$

Applying monotonicity of the function $\kappa \mapsto \tilde\measuredangle^{\kappa}(p\,{}^{q}_{r})$ (1.1*d*) shows that

$$\tilde\measuredangle^{\kappa}(p\,{}^{q}_{r}) + \tilde\measuredangle^{\kappa}(p\,{}^{r}_{s}) + \tilde\measuredangle^{\kappa}(p\,{}^{s}_{q}) \leqslant 2\cdot\pi.$$

Since the quadruple p,q,r,s is arbitrary, the if part follows.

Now let us prove the only-if part. Denote by σ the exact upper bound on values $K \geqslant \kappa$ such that all model triangles with the vertices p,q,r,s are defined.

Recall that $\tilde{\measuredangle}^{K+}(p\,{}^q_r)$ denotes extended angle (8.48). Observe that if

❸ $$\tilde{\measuredangle}^{K+}(p\,{}^q_r) + \tilde{\measuredangle}^{K+}(p\,{}^r_s) + \tilde{\measuredangle}^{K+}(p\,{}^s_q) = 2\cdot\pi$$

for some $\sigma \geqslant K \geqslant \kappa$, then the quadruple admits a distance-preserving embedding into $\mathbb{M}^2(K)$.

Observe that the left-hand side of ❸ is continuous in K. Since $\mathcal{L}$ is CBB(κ), for $K = \kappa$ the left-hand side cannot exceed $2\cdot\pi$. Therefore it remains smaller than $2\cdot\pi$ for all $\sigma \geqslant K \geqslant \kappa$; moreover the same holds for all permutations of the labels p,q,r,s.

Note that we can assume the perimeter of the triple q,r,s is $2\cdot\varpi^\sigma$, and use this and the overlap lemma (10.2) to arrive at a contradiction.

According to our definition, the real line is CBB(κ) for any $\kappa \in \mathbb{R}$, but it does not satisfy the property for $\kappa > 0$. The condition $\kappa \leqslant 0$ was used just once to ensure that the κ-model triangles with the vertices p,q,r,s are defined. One can assume instead that perimeters of all triangles in $\mathcal{L}$ are at most $2\cdot\varpi^\kappa$. This condition holds for most complete length CBB(κ) spaces of dimension at least 2.

10.9. Let $\tilde{p},\tilde{x}_1,\dots,\tilde{x}_n$ be the array in $\mathbb{E}^n$ provided by the $(1{+}n)$-point comparison (10.8). We may assume that $\tilde{p}$ is the origin of $\mathbb{E}^n$.

Consider an $n{\times}n$-matrix $\tilde{M}$ with components

$$\tilde{m}_{i,j} = \tfrac{1}{2}\cdot(|\tilde{x}_i - \tilde{p}|^2 + |\tilde{x}_j - \tilde{p}|^2 - |\tilde{x}_i - \tilde{x}_j|^2).$$

Note that $\tilde{m}_{i,j} = \langle \tilde{x}_i, \tilde{x}_j\rangle$. Therefore $\tilde{M} = A\cdot A^\top$ for an $n{\times}n$-matrix A that defines a linear transformation sending the standard basis to the array $\tilde{x}_1,\dots,\tilde{x}_n$. Therefore

$$\boldsymbol{s}\cdot\tilde{M}\cdot\boldsymbol{s}^\top = |A^\top\cdot\boldsymbol{s}^\top|^2 \geqslant 0$$

for any vector $\boldsymbol{s}$. Further show that

$$\boldsymbol{s}\cdot M\cdot\boldsymbol{s}^\top \geqslant \boldsymbol{s}\cdot\tilde{M}\cdot\boldsymbol{s}^\top$$

for any vector $\boldsymbol{s} = (s_1,\dots,s_n)$ with nonnegative components.

10.11. It is sufficient to construct a metric on the set of points $\{p, x^1, x^2, x^3, x^4\}$ that does not satisfy $(4{+}1)$-comparison but does satisfy all $(3{+}1)$-comparisons. To do this, set x^i to be the vertices of a regular tetrahedron in $\mathbb{E}^3$. Suppose p is its center and reduce the distances $|p - x^i|$ slightly.

10.12. By the $(1+n)$-point comparison (10.8), there is a point array $\tilde p, \tilde a^0, \dots, \tilde a^m \in \mathbb{M}^{m+1}(\kappa)$ such that

$$|\tilde p - \tilde a^i| = |p - a^i| \quad \text{and} \quad |\tilde a^i - \tilde a^j| \geqslant |a^i - a^j|$$

for all i and j.

For each i, set $\tilde\xi^i = \uparrow_{[\tilde p \tilde a^i]} \in \mathbb{S}^m = \Sigma_{\tilde p}(\mathbb{M}^{m+1}(\kappa))$. Note that

$$|\tilde\xi^i - \tilde\xi^j|_{\mathbb{S}^m} \geqslant \measuredangle^\kappa(p\,{}^{a^i}_{a^j}) > \tfrac{\pi}{2}.$$

Consider two matrices S and $\tilde S$ with components $s_{i,j} = \langle \tilde\xi^i, \xi^j\rangle$ and $\tilde s_{i,j} = \cos[\measuredangle^\kappa(p\,{}^{a^i}_{a^j})]$. By construction, $S \geqslant 0$.

Note that $s_{i,j} \leqslant \tilde s_{i,j} \leqslant 0$ if $i \ne j$ and $s_{i,j} = \tilde s_{i,j} = 1$ if $i = j$. Therefore $\tilde S \geqslant 0$. The latter implies we can assume

$$|\tilde\xi^i - \tilde\xi^j|_{\mathbb{S}^m} = \measuredangle^\kappa(p\,{}^{a^i}_{a^j})$$

for each i and j. Whence the statement follows.

10.18. Set $\tilde Q = \mathrm{Conv}\{\tilde x^0, \tilde x^1, \dots, \tilde x^{\hbar}\}$. By Kirszbraun's theorem, the map $\tilde x^i \mapsto x^i$ can be extended to a short map $F\colon \tilde Q \to \mathcal{L}$; it remains to show that the map F is distance-preserving.

Consider the gradient logarithm map $\mathrm{lg}_{x_0}\colon \mathcal{L} \to \mathrm{T}_{x_0}$ corresponding to the gradient exponential map gexp_{x_0} (see Section 16G). The map lg_{x_0} also is short. Observe that the composition $\mathrm{lg}_{x_0} \circ F$ is distance-preserving. Therefore F is distance-preserving; in particular we can take $Q = F(\tilde Q)$.

10.19. Consider vectors $v^i = \mathrm{lg}_{x^0} x^i \in \mathrm{T}_{x^0}$. Show that all the v^i lie in a linear subspace of T_{x^0} and that $x^i \mapsto v^i$ is distance-preserving. It follows that we can identify the convex hull K of the v^i with the convex hull of the $\tilde x^i$.

Note that the gradient exponential map gexp_{x_0} maps v^i to x^i. By assumption,

❹ $$|v^i - v^j| = |x^i - x^j|$$

for all i and j. By 16.35, gexp_{x_0} is a short map. By ❹, gexp_{x_0} cannot be strictly short at a pair of points in K. That is, gexp_{x_0} is distance-preserving on K.

10.20. Apply 10.17 for each of the following maps

- ◇ $f_0\colon \tilde x \mapsto x$, $\tilde p^1 \mapsto p^1$, $\tilde q^1 \mapsto q^1$;
- ◇ $f_i\colon \tilde p^i \mapsto p^i$, $\tilde p^{i+1} \mapsto p^{i+1}$, $\tilde q^i \mapsto q^i$, $\tilde q^{i+1} \mapsto q^{i+1}$ for $1 \leqslant i < n$;
- ◇ $f_n\colon \tilde y \mapsto y$, $\tilde p^n \mapsto p^n$, $\tilde q^n \mapsto q^n$.

Denote by F_i the short extension of f_i. Observe that $F_{i-1}(\tilde z_i) = F_i(\tilde z_i)$ for each i and use it.

10.23. Consider the space $\mathcal{Y}^{\mathcal{X}}$ of all maps $\mathcal{X} \to \mathcal{Y}$ equipped with the product topology.

Denote by $\mathfrak{S}_F$ the set of maps $h \in \mathcal{Y}^{\mathcal{X}}$ such that the restriction $h|_F$ is short and agrees with f in $F \cap A$. Note that the sets $\mathfrak{S}_F \subset \mathcal{Y}^{\mathcal{X}}$ are closed and any finite intersection of these sets is nonempty.

According to Tikhonov's theorem, $\mathcal{Y}^{\mathcal{X}}$ is compact. By the finite intersection property, the intersection $\bigcap_F \mathfrak{S}_F$ for all finite sets $F \subset X$ is nonempty. Hence the statement follows.

10.24. The Kuratowsky embedding is a distance-preserving map of $\mathcal{X}$ into the space of bounded functions $\mathcal{X}$ equipped with the metric induced by the sup-norm (Section 2G). It remains to show that the latter space is injective.

The second part of the exercise is a classical result of John Isbell [70] which was rediscovered several times after him.

10.25. To prove the only-if part it is sufficient to consider only the case of two-point spaces $\mathcal{Z}$.

It remains to prove the if part.

By Zorn's lemma, we may assume that the short map $f\colon Q \to \mathcal{Z}$ is defined on a maximal subset $Q \subset \mathcal{X}$; that is, f cannot be extended to a larger set as a short map. Note that in this case Q is a closed subset of $\mathcal{X}$.

Assume there is a point $x \in \mathcal{X}\backslash Q$. Denote by $\bar x$ a point in Q that lies at minimal distance from x. By the ultratriangle inequality, $|\bar x - y| \leqslant |x - y|$ for any $y \in Q$. Assigning $f(x) = f(\bar x)$, we extend f to $Q \cup \{x\}$ as a short map — a contradiction. Therefore $Q = \mathcal{X}$, hence the statement.

11.4. It is sufficient to show that the natural map $\mathcal{B} \times_g \mathcal{F} \to \mathcal{B} \times_f \mathcal{F}$ is short. The latter follows from the fiber-independence theorem (11.3).

11.5. Show and apply that any geodesic path in $\mathrm{Cone}^\kappa \mathcal{F}$ projects to a geodesic in $\mathcal{F}$ of length less than π.

11.9. By 11.6a, the space $\mathcal{U}$, $\mathcal{V}$, or $\mathcal{U} \star \mathcal{V}$ is CBB(1) if and only if $\mathrm{Cone}\,\mathcal{U}$, $\mathrm{Cone}\,\mathcal{V}$, or $\mathrm{Cone}(\mathcal{U} \star \mathcal{V}) = \mathrm{Cone}\,\mathcal{U} \times \mathrm{Cone}\,\mathcal{V}$ is CBB(0) respectively.

By 11.7a, the space $\mathcal{U}$, $\mathcal{V}$, or $\mathcal{U} \star \mathcal{V}$ is CAT(1) if and only if $\mathrm{Cone}\,\mathcal{U}$, $\mathrm{Cone}\,\mathcal{V}$, or $\mathrm{Cone}(\mathcal{U} \star \mathcal{V}) = \mathrm{Cone}\,\mathcal{U} \times \mathrm{Cone}\,\mathcal{V}$ is CAT(0) respectively.

It remains to show that the product of two spaces is CBB(0) or CAT(0) if and only if each space is CBB(0) or CAT(0) respectively.

12.4. Assume $\mathcal{P}$ is not CAT(0). Then by 12.2, a link Σ of some simplex contains a closed local geodesic α with length $4{\cdot}\ell < 2{\cdot}\pi$. We can assume that Σ has minimal possible dimension; then by 12.2, 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, and parametrize them by $[-\ell,\ell]$. Fix a small $\delta>0$ and consider the two curves in $\operatorname{Cone}\Sigma$ given in polar coordinates by

$$\gamma_i(t)=(\alpha_i(\arctan\tfrac{t}{\delta}),\sqrt{\delta^2+t^2}).$$

Observe that the curves γ_1 and γ_2 are geodesics in $\operatorname{Cone}\Sigma$ having common endpoints.

Observe that a small neighborhood of the tip of $\operatorname{Cone}\Sigma$ admits a distance-preserving 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 where α_1 (and therefore α_2) is not length-minimizing.

Pass to a maximal length-minimizing arc $\bar\alpha_1$ of α_1. Since Σ is locally CAT(1), by the no-conjugate-point theorem (9.44) 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 $\bar\alpha_1$, $\bar\alpha_2$.

Remark. By 9.8 the converse holds as well.

12.6. Apply 13.1, 12.5, and 12.2.

12.8. Observe and use that (1) in the barycentric subdivision every vertex corresponds to a simplex of the original triangulation, and (2) a simplex of the subdivision corresponds to a decreasing sequence of simplexes in the original triangulation.

12.13. 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 (12.11). The only difference is that a geodesic may spend time at least π on each visit to Star_v.

Remark. Note that it is not sufficient to assume only that all the dihedral angles of the simplexes are at least $\frac{\pi}{2}$. Indeed, the two-dimensional sphere with the interior of a small rhombus removed is a spherical polyhedral space glued from four triangles with angles at least $\frac{\pi}{2}$. On the other hand, the boundary of the rhombus is a closed local geodesic in this space and has length less than 2π. Therefore the space cannot be CAT(1).

12.14. Observe that if we glue two copies of spaces along A_i, then the copies of A_j for some $j\neq i$ form a convex subset in the glued space. Use this and the Reshetnyak gluing theorem (9.38) n times, once for each label of the edges.

12.15. The space $\mathcal{T}_n$ has a natural cone structure whose vertex is 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 by right-angled spherical simplexes. Each simplex corresponds to the combinatorics of a possibly degenerate tree.

Note that the link of any simplex of this triangulation satisfies the no-triangle condition. Indeed, fix a simplex $\triangle$ of the complex; suppose it is described by a possibly degenerate topological tree t. A triangle in the link of $\triangle$ can be described by three ways to resolve a degeneracy of t by adding one edge, where (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 12.9 our complex is flag. It remains to apply the flag condition (12.11) and 11.6*a*.

12.16. Apply the flag condition (12.11) and Theorem 11.7*a*.

13.8. Consider a cube in the Hilbert space of all sequences $x_1, x_2, \dots$ such that $|x_i| \leqslant 1$ and $x_1^2 + x_2^2 + \dots < \infty$ with the metric induced by the ℓ^2-norm.

13.13 and 13.14. Apply the strong angle lemmas 9.34 and 8.42.

13.25. Since α is Lipschitz, so is $f \circ \alpha$. By the standard Rademacher theorem, the derivative $(f\circ\alpha)'$ is defined almost everywhere. In particular,

$$(\boldsymbol{d}_{\alpha(t)} f)(\alpha^+(t)) + (\boldsymbol{d}_{\alpha(t)} f)(\alpha^-(t)) \overset{a.e.}{=\!=} 0.$$

Further, by the extended Rademacher theorem (more precisely its 1-dimensional case; see Proposition 13.9), we have

$$\alpha^+(t) + \alpha^-(t) \overset{a.e.}{=\!=} 0.$$

In particular,

$$\langle \nabla_{\alpha(t)} f, \alpha^+(t)\rangle + \langle \nabla_{\alpha(t)} f, \alpha^-(t)\rangle \overset{a.e.}{=\!=} 0.$$

Finally, by the definition of gradient, we have

$$\langle \nabla_{\alpha(t)} f, \alpha^\pm(t)\rangle \geqslant (\boldsymbol{d}_{\alpha(t)} f)(\alpha^\pm(t)).$$

Hence the result follows.

13.32. Without loss of generality, we may assume that geodesics $[pa]$ and $[pb]$ are uniquely defined. Applying 13.14, we have

$$\begin{aligned}(\boldsymbol{d}_p \operatorname{dist}_a)(\nabla_p \operatorname{dist}_b) &= -\langle \uparrow_{[pa]}, \nabla_p \operatorname{dist}_b \rangle \leqslant \\ &\leqslant -\boldsymbol{d}_p \operatorname{dist}_b(\uparrow_{[pa]}) = \\ &= \langle \uparrow_{[pb]}, \uparrow_{[pa]} \rangle = \\ &= \cos \measuredangle[p\,{}^a_b] \leqslant \\ &\leqslant \cos \tilde\measuredangle^{\kappa}(p\,{}^a_b).\end{aligned}$$

15.17. Suppose that $\mathcal{L}$ is infinite-dimensional. Denote by $\Omega_m \subset \mathcal{L}$ the set of all points p with $\operatorname{rank}_p \geqslant m$. Evidently $\Omega_1 \supset \Omega_2 \supset \dots$, and Ω_m is open for each m.

By 15.6C, each Ω_m is dense in $\mathcal{L}$. Hence there is a G-delta dense set of points $p \in \mathcal{L}$ such that $\operatorname{rank}_p = \infty$. It follows that Σ_p is not compact.

16.1. Choose a finite sequence $t_0 < \dots < t_n$. Denote by Φ_{t_i} the composition of projections to $K_{t_0}, \dots, K_{t_i}$. Pass to a limit of the Φ_{t_i} as the sequence becomes denser in the parameter interval. Observe that the limit φ_t does not depend on the choice of the sequences. The exercise follows.

16.8. Let $\ell(t) = |\alpha(t) - \alpha(t_3)|$. Note that

$$\ell'(t) \leqslant -\langle \nabla_{\alpha(t)} f, \uparrow_{[\alpha(t)\alpha(t_3)]} \rangle.$$

Observe that the function $t \mapsto f \circ \alpha(t)$ is nondecreasing; in particular, $f(\alpha(t_1)) \leqslant f(\alpha(t_2)) \leqslant f(\alpha(t_3))$. Therefore

$$\langle \nabla_{\alpha(t)} f, \uparrow_{[\alpha(t)\alpha(t_3)]} \rangle \geqslant \boldsymbol{d}_{\alpha(t)} f(\uparrow_{[\alpha(t)\alpha(t_3)]}) \geqslant 0$$

for any $t \in [t_1, t_2]$. Therefore $\ell' \leqslant 0$ for any $t \in [t_1, t_2]$. Hence the statement.

16.10. Without loss of generality, we may assume that $(f \circ \alpha)'(t) > 0$ for any t.

Let $\hat{\alpha}$ be the arclength reparametrization of α. Note that

$$(f \circ \hat{\alpha})'(s) \geqslant |\nabla_{\hat{\alpha}(s)} f|$$

almost everywhere. Therefore, by Theorem 16.3, $\hat{\alpha}$ is a gradient-like curve. It remains to apply Lemma 16.9.

16.11. Use 16.10 to prove the only-if part.

To prove the if part, set $h(z) = \frac{1}{2} \cdot |x - z|^2$. If α is an f-gradient curve, then

$$\begin{aligned}(h \circ \alpha)^+ &\geqslant |\alpha(t) - x| \cdot \langle \uparrow_{[\alpha(t)x]}, \nabla_{\alpha(t)} f \rangle \geqslant \\ &\geqslant |\alpha(t) - x| \cdot \boldsymbol{d}_{\alpha(t)} f(\uparrow_{[\alpha(t)x]}) \geqslant \\ &\geqslant f(x) - f \circ \alpha(t).\end{aligned}$$

It remains to integrate the inequality and observe that $f \circ \alpha$ is nondecreasing.

16.34. Consider (x, κ)- and (z, κ)-radial curves that start at y and observe that they form a geodesic from x to z.

16.37. Set $q = p + v$ and $q' = \operatorname{gexp}_p v$. By radial comparison, $|q' - x| \leqslant |q - x|$ for any $x \in \mathcal{L}$. If $q \in \mathcal{L}$, this implies that $q = q'$. Otherwise note that q' lies on the boundary line of $\mathcal{L}$, and $\operatorname{proj}(q)$ is the only point on this line that satisfies the inequality.

16.38. By the angle comparison, $|\nabla_x \operatorname{dist}_p| \geqslant -\cos\tilde{\measuredangle}^\kappa(x\,{}^p_q)$.

Choose a (p,κ)-radial curve α that starts at p. Observe that

$$(\operatorname{dist}_p\circ\alpha)^+(t)\geqslant -|\alpha^+(t)|\cdot\cos\tilde{\measuredangle}^\kappa(\alpha(t)\,{}^p_q)$$

and

$$(\operatorname{dist}_q\circ\alpha)^+(t)\geqslant -|\alpha^+(t)|.$$

Therefore $t\mapsto\tilde{\measuredangle}^\kappa(q\,{}^{\alpha(t)}_{p})$ is nondecreasing, hence the result.

16.39. Choose a regular point $p\in\mathcal{L}$. Observe that $\operatorname{gexp}^\kappa_p(\overline{\mathrm{B}}[0,R]_{\mathrm{T}_p}) = {}$ $= \mathcal{L}$.

Use 16.38 to show that $\operatorname{gexp}^\kappa_p(\partial\overline{\mathrm{B}}[0,R]_{\mathrm{T}_p})\supset\partial\mathcal{L}$ and apply 16.35*b*.

Index

Bibliography

[1] A. D. Aleksandrov and V. A. Zalgaller. *Intrinsic geometry of surfaces*. 1967. Translated from the Russian by J. M. Danskin. Translations of Mathematical Monographs, Vol. 15.

[2] S. Alexander and R. Bishop. "Comparison theorems for curves of bounded geodesic curvature in metric spaces of curvature bounded above". *Differential Geom. Appl.* 6.1 (1996), 67–86.

[3] S. Alexander and R. Bishop. "The Hadamard-Cartan theorem in locally convex metric spaces". *Enseign. Math. (2)* 36.3-4 (1990), 309–320.

[4] S. Alexander and R. Bishop. "Warped products admitting a curvature bound". *Adv. Math.* 303 (2016), 88–122.

[5] S. Alexander and R. Bishop. "Curvature bounds for warped products of metric spaces". *Geom. Funct. Anal.* 14.6 (2004), 1143–1181.

[6] S. Alexander, V. Kapovitch, and A. Petrunin. "Alexandrov meets Kirszbraun". *Proceedings of the Gökova Geometry-Topology Conference 2010*. 2011, 88–109.

[7] S. Alexander, V. Kapovitch, and A. Petrunin. "An optimal lower curvature bound for convex hypersurfaces in Riemannian manifolds". *Illinois J. Math.* 52.3 (2008), 1031–1033.

[8] V. Kapovitch S. Alexander and A. Petrunin. *An invitation to Alexandrov geometry: CAT(0) spaces*. SpringerBriefs in Mathematics. 2019.

[9] А. Д. Александров. *Внутренняя геометрия выпуклых поверхностей*. 1948.

[10] A. Alexandroff. "The inner geometry of an arbitrary convex surface". *C. R. (Doklady) Acad. Sci. URSS (N.S.)* 32 (1941), 467–470.

[11] A. D. Alexandrow. "Über eine Verallgemeinerung der Riemannschen Geometrie". *Schr. Forschungsinst. Math.* 1 (1957), 33–84.

[12] А. Д. Александров. «Одна теорема о треугольниках в метрическом пространстве и некоторые ее приложения». *Тр. МИАН СССР* 38 (1951), 5—23.

[13] L. Ambrosio, N. Gigli, and G. Savaré. *Gradient flows in metric spaces and in the space of probability measures*. Second. Lectures in Mathematics ETH Zürich. 2008.

[14] W. Ballmann. "Singular spaces of nonpositive curvature". *Sur les groupes hyperboliques d'après Mikhael Gromov (Bern, 1988)*. Vol. 83. Progr. Math. 1990, 189–201.

[15] W. Ballmann. *Lectures on spaces of nonpositive curvature*. Vol. 25. DMV Seminar. With an appendix by Misha Brin. 1995.

[16] W. Ballmann. *On the geometry of metric spaces*. MPIM. 2004. URL: `https://people.mpim-bonn.mpg.de/hwbllmnn/archiv/sin40827.pdf`.

[17] В. Н. Берестовский. «Задача Борсука о метризации полиэдра». *Докл. АН СССР* 268.2 (1983), 273—277.

[18] V. N. Berestovskij and I. G. Nikolaev. "Multidimensional generalized Riemannian spaces". *Geometry, IV*. Vol. 70. Encyclopaedia Math. Sci. 1993, 165–243, 245–250.

[19] I. D. Berg and I. G. Nikolaev. “Quasilinearization and curvature of Aleksandrov spaces”. *Geom. Dedicata* 133 (2008), 195–218.

[20] A. S. Besicovitch and P. A. P. Moran. “The measure of product and cylinder sets”. *J. London Math. Soc.* 20 (1945), 110–120.

[21] L. Billera, S. Holmes, and K. Vogtmann. “Geometry of the space of phylogenetic trees”. *Adv. in Appl. Math.* 27.4 (2001), 733–767.

[22] W. Blaschke. *Kreis und Kugel*. 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. 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 [Fundamental Principles of Mathematical Sciences]. 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.

[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. “Uniform estimates on the number of collisions in semi-dispersing billiards”. *Ann. of Math. (2)* 147.3 (1998), 695–708.

[30] D. Burago, S. Ferleger, and A. Kononenko. “Topological entropy of semi-dispersing billiards”. *Ergodic Theory Dynam. Systems* 18.4 (1998), 791–805.

[31] D. Burago, S. Ferleger, and A. Kononenko. “Unfoldings and global bounds on the number of collisions for generalized semi-dispersing billiards”. *Asian J. Math.* 2.1 (1998), 141–152.

[32] D. Burago, S. Ferleger, and A. Kononenko. “A geometric approach to semi-dispersing billiards”. *Hard ball systems and the Lorentz gas*. Vol. 101. Encyclopaedia Math. Sci. 2000, 9–27.

[33] D. Burago, S. Ferleger, B. Kleiner, and A. Kononenko. “Gluing copies of a 3-dimensional polyhedron to obtain a closed nonpositively curved pseudomanifold”. *Proc. Amer. Math. Soc.* 129.5 (2001), 1493–1498.

[34] Yu. Burago, M. Gromov, and G. Perelman. “A. D. Aleksandrov spaces with curvatures bounded below”. *Russian Math. Surveys* 47.2 (1992), 1–58.

[35] H. Busemann. “Spaces with non-positive curvature”. *Acta Math.* 80 (1948), 259–310.

[36] S. V. Bujalo. “Shortest paths on convex hypersurfaces of a Riemannian space”. *Zap. Naučn. Sem. Leningrad. Otdel. Mat. Inst. Steklov. (LOMI)* 66 (1976). Studies in topology, II, 114–132, 207.

[37] É. Cartan. *Leçons sur la Géométrie des Espaces de Riemann*. 1928.

[38] R. Charney and M. Davis. “Singular metrics of nonpositive curvature on branched covers of Riemannian manifolds”. *Amer. J. Math.* 115.5 (1993), 929–1009.

[39] R. Charney and M. Davis. “Strict hyperbolization”. *Topology* 34.2 (1995), 329–350.

[40] J. Cheeger and T. Colding. “Lower bounds on Ricci curvature and the almost rigidity of warped products”. *Ann. of Math. (2)* 144.1 (1996), 189–237.

[41] J. Cheeger. *Comparison and finiteness theorems for riemannian manifolds*. Thesis (Ph.D.)–Princeton University. 1967.

[42] J. Cheeger and D. Gromoll. “The splitting theorem for manifolds of nonnegative Ricci curvature”. *J. Differential Geometry* 6 (1971/72), 119–128.

[43] S. Cohn-Vossen. „Totalkrümmung und geodätische Linien auf einfachzusammenhängenden offenen vollständigen Flächenstücken“. *Mat. Sb.* 1(43).2 (1936).

[44] A. Daniilidis, O. Ley, and S. Sabourau. "Asymptotic behaviour of self-contracted planar curves and gradient orbits of convex functions". *J. Math. Pures Appl. (9)* 94.2 (2010), 183–199.

[45] M. Davis. "Exotic aspherical manifolds". *Topology of high-dimensional manifolds, No. 1, 2 (Trieste, 2001)*. Vol. 9. ICTP Lect. Notes. 2002, 371–404.

[46] M. Davis. "Groups generated by reflections and aspherical manifolds not covered by Euclidean space". *Ann. of Math. (2)* 117.2 (1983), 293–324.

[47] M. Davis. "Exotic aspherical manifolds". *Topology of high-dimensional manifolds, No. 1, 2 (Trieste, 2001)*. Vol. 9. ICTP Lect. Notes. 2002, 371–404.

[48] D. Edwards. "The structure of superspace" (1975), 121–133.

[49] R. Engelking. *Dimension theory*. 1978.

[50] J.-H. Eschenburg. "The splitting theorem for space-times with strong energy condition". *J. Differential Geom.* 27.3 (1988), 477–491.

[51] L. Ferreira and J. C. Valencia-Guevara. "Gradient flows of time-dependent functionals in metric spaces and applications to PDEs". *Monatsh. Math.* 185.2 (2018), 231–268.

[52] Z. Frolık. "Concerning topological convergence of sets". *Czechoslovak Math. J* 10(85) (1960), 168–180.

[53] N. Gigli. "An overview of the proof of the splitting theorem in spaces with non-negative Ricci curvature". *Anal. Geom. Metr. Spaces* 2.1 (2014), 169–213.

[54] N. Gigli. *The splitting theorem in non-smooth context*. 2013. arXiv: 1302.5555.

[55] R. E. Greene and H. Wu. "On the subharmonicity and plurisubharmonicity of geodesically convex functions". *Indiana Univ. Math. J.* 22 (1972/73), 641–653.

[56] M. Gromov. *Metric structures for Riemannian and non-Riemannian spaces*. English. Modern Birkhäuser Classics. Based on the 1981 French original, With appendices by M. Katz, P. Pansu and S. Semmes, Translated from the French by Sean Michael Bates. 2007.

[57] M. Gromov. "Groups of polynomial growth and expanding maps". *Inst. Hautes Études Sci. Publ. Math.* 53 (1981), 53–73.

[58] M. Gromov. "Almost flat manifolds". *J. Differential Geom.* 13.2 (1978), 231–241.

[59] 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.

[60] M. Gromov. "Curvature, diameter and Betti numbers". *Comment. Math. Helv.* 56.2 (1981), 179–195.

[61] M. Gromov. "Hyperbolic groups". *Essays in group theory*. Vol. 8. Math. Sci. Res. Inst. Publ. 1987, 75–263.

[62] K. Grove and P. Petersen. "Bounding homotopy types by geometry". *Ann. of Math. (2)* 128.1 (1988), 195–206.

[63] K. Grove and B. Wilking. "A knot characterization and 1-connected nonnegatively curved 4-manifolds with circle symmetry". *Geom. Topol.* 18.5 (2014), 3091–3110.

[64] 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.

[65] S. Halbeisen. "On tangent cones of Alexandrov spaces with curvature bounded below". *Manuscripta Math.* 103.2 (2000), 169–182.

[66] F. Hausdorff. *Grundzüge der Mengenlehre*. 1914.

[67] J. Heinonen, P. Koskela, N. Shanmugalingam, and J. Tyson. *Sobolev spaces on metric measure spaces*. Vol. 27. New Mathematical Monographs. An approach based on upper gradients. 2015.

[68] W.-Y. Hsiang and B. Kleiner. "On the topology of positively curved 4-manifolds with symmetry". *J. Differential Geom.* 29.3 (1989), 615–621.

[69] W. Hurewicz and H. Wallman. *Dimension Theory*. Princeton Mathematical Series, v. 4. 1941.

[70] J. R. Isbell. “Six theorems about injective metric spaces”. *Comment. Math. Helv.* 39 (1964), 65–76.

[71] S. Ivanov. *To what extent is convexity a local property?* MathOverflow. (version: 2012-06-09). eprint: `https://mathoverflow.net/q/22091`.

[72] J. Jost. “Nonlinear Dirichlet forms”. *New directions in Dirichlet forms.* Vol. 8. AMS/IP Stud. Adv. Math. 1998, 1–47.

[73] J. Jost. *Nonpositive curvature: geometric and analytic aspects.* Lectures in Mathematics ETH Zürich. 1997.

[74] C. Jun. “Pursuit-evasion and time-dependent gradient flow in singular spaces”. Thesis (Ph.D.)–University of Illinois at Urbana-Champaign. PhD thesis. 2012.

[75] C. Jun. *Time-dependent gradient curves on CAT(0) spaces.* arXiv: `1801.04583 [math.MG]`.

[76] B. Kirchheim. “Rectifiable metric spaces: local structure and regularity of the Hausdorff measure”. *Proc. Amer. Math. Soc.* 121.1 (1994), 113–123.

[77] B. Kleiner. “The local structure of length spaces with curvature bounded above”. *Math. Z.* 231.3 (1999), 409–456.

[78] B. Kleiner and B. Leeb. “Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings”. *Inst. Hautes Études Sci. Publ. Math.* 86 (1997), 115–197 (1998).

[79] U. Lang and V. Schroeder. “Kirszbraun’s theorem and metric spaces of bounded curvature”. *Geom. Funct. Anal.* 7.3 (1997), 535–560.

[80] U. Lang and V. Schroeder. “On Toponogov’s comparison theorem for Alexandrov spaces”. *Enseign. Math.* 59.3-4 (2013), 325–336.

[81] N. Lebedeva and A. Petrunin. “Curvature bounded below: a definition a la Berg–Nikolaev”. *Electron. Res. Announc. Math. Sci.* 17 (2010), 122–124.

[82] N. Lebedeva and A. Petrunin. “Local characterization of polyhedral spaces”. *Geom. Dedicata* 179 (2015), 161–168.

[83] N. Lebedeva, A. Petrunin, and V. Zolotov. “Bipolar comparison”. *Geom. Funct. Anal.* 29.1 (2019), 258–282.

[84] N. Lebedeva and A. Petrunin. *5-point Toponogov theorem.* URL: `https://github.com/anton-petrunin/toyoda`.

[85] N. Lebedeva and A. Petrunin. “5-point CAT(0) spaces after Tetsu Toyoda”. *Anal. Geom. Metr. Spaces* 9.1 (2021), 160–166.

[86] И. М. Либерман. «Геодезические линии на выпуклых поверхностях». *Докл. АН СССР* 32.2 (1941), 310—312.

[87] A. Lytchak and A. Petrunin. *Short retractions of CAT(1) spaces.* 2020. arXiv: `2002.09762 [math.DG]`.

[88] A. Lytchak and A. Petrunin. *Short retractions of CAT(1) spaces.* 2020. arXiv: `2002.09762 [math.DG]`.

[89] A. Lytchak. “Differentiation in metric spaces”. *St. Petersburg Math. J.* 16.6 (2005), 1017–1041.

[90] A. Lytchak. “Open map theorem for metric spaces”. *St. Petersburg Math. J.* 17.3 (2006), 477–491.

[91] 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.

[92] G. A. Margulis and G. D. Mostow. “Some remarks on the definition of tangent cones in a Carnot-Carathéodory space”. *J. Anal. Math.* 80 (2000), 299–317.

[93] U. Mayer. “Gradient flows on nonpositively curved metric spaces and harmonic maps”. *Comm. Anal. Geom.* 6.2 (1998), 199–253.

[94] A. Mielke, R. Rossi, and G. Savaré. “Nonsmooth analysis of doubly nonlinear evolution equations”. *Calc. Var. Partial Differential Equations* 46.1-2 (2013), 253–310.

[95] А. Д. Милка. «Метрическое строение одного класса пространств, содержащих прямые линии.». *Укр. геометр, сб.* 4 (1967), 43—48.

[96] N. Monod. "Superrigidity for irreducible lattices and geometric splitting". *J. Amer. Math. Soc.* 19.4 (2006), 781–814.

[97] J. Nagata. *Modern dimension theory*. Revised. Vol. 2. Sigma Series in Pure Mathematics. 1983.

[98] F. Nazarov. *Intrinsic metric with no geodesics*. MathOverflow. (version: 2010-02-18). eprint: `http://mathoverflow.net/q/15720`.

[99] S. Ohta. "Gradient flows on Wasserstein spaces over compact Alexandrov spaces". *Amer. J. Math.* 131.2 (2009), 475–516.

[100] D. Panov and A. Petrunin. "Ramification conjecture and Hirzebruch's property of line arrangements". *Compos. Math.* 152.12 (2016), 2443–2460.

[101] G. Perelman and A. Petrunin. *Quasigeodesics and Gradient curves in Alexandrov spaces*. eprint: `http://www.math.psu.edu/petrunin/`.

[102] G. Ya. Perelman. "Elements of Morse theory on Aleksandrov spaces". *St. Petersburg Math. J.* 5.1 (1994), 205–213.

[103] G. Perelman. *Alexandrov paces with curvature bounded from below II.* Preprint LOMI, 1991. 35pp. eprint: `http://www.math.psu.edu/petrunin/`.

[104] Anton Petrunin and Stephan Stadler. "Metric-minimizing surfaces revisited". *Geom. Topol.* 23.6 (2019), 3111–3139.

[105] A. Petrunin. "A globalization for non-complete but geodesic spaces". *Math. Ann.* 366.1-2 (2016), 387–393.

[106] A. Petrunin. "Alexandrov meets Lott-Villani-Sturm". *Münster J. Math.* 4 (2011), 53–64.

[107] A. Petrunin. "Semiconcave functions in Alexandrov's geometry". *Surveys in differential geometry. Vol. XI.* Vol. 11. Surv. Differ. Geom. 2007, 137–201.

[108] R. Phelps. *Convex functions, monotone operators and differentiability.* Second. Vol. 1364. Lecture Notes in Mathematics. 1993.

[109] P. Pizzetti. "Paragone fra due triangoli a lati uguali". *Atti della Reale Accademia dei Lincei, Rendiconti (5). Classe di Scienze Fisiche, Matematiche e Naturali* 16.1 (1907), 6–11.

[110] Conrad Plaut. *Toponogov's theorem for metric spaces.* 1990.

[111] C. Plaut. "Spaces of Wald curvature bounded below". *J. Geom. Anal.* 6.1 (1996), 113–134.

[112] C. Plaut. "Metric spaces of curvature $\geq k$". *Handbook of geometric topology.* 2002, 819–898.

[113] L. Pontrjagin. "Sur une hypothèse fondamentale de la théorie de la dimension." *C. R. Acad. Sci., Paris* 190 (1930), 1105–1107.

[114] Ю. Г. Решетняк. «Двумерные многообразия ограниченной кривизны». *Геометрия – 4.* Т. 70. Итоги науки и техн. Сер. Соврем. пробл. мат. Фундам. направления. 1989, 7—189.

[115] Yu. G. Reshetnyak. "Inextensible mappings in a space of curvature no greater than K". *Siberian Mathematical Journal* 9.4 (1968), 683–689.

[116] W. Rinow. *Die innere Geometrie der metrischen Räume.* Die Grundlehren der mathematischen Wissenschaften, Bd. 105. 1961.

[117] W. Rudin. "Homogeneity problems in the theory of Čech compactifications". *Duke Math. J.* 23 (1956), 409–419.

[118] I. Kh. Sabitov. "Around the proof of the Legendre–Cauchy lemma on convex polygons". *Siberian Math. J.* 45.4 (2004), 740–762.

[119] T. Sato. "An alternative proof of Berg and Nikolaev's characterization of CAT(0)-spaces via quadrilateral inequality". *Arch. Math. (Basel)* 93.5 (2009), 487–490.

[120] T. Foertsch and V. Schroeder. "Minkowski versus Euclidean rank for products of metric spaces". *Adv. Geom.* 2.2 (2002), 123–131.

[121] G. Savaré. "Gradient flows and diffusion semigroups in metric spaces under lower curvature bounds". *C. R. Math. Acad. Sci. Paris* 345.3 (2007), 151–154.

[122] В. А. Шарафутдинов. «Теорема Погорелова — Клингенберга для многообразий, гомеоморфных $\mathbb{R}^n$». *Сибирский математический журнал* 18.4 (1977), 915—925.

[123] K. Shiohama. *An introduction to the geometry of Alexandrov spaces*. Vol. 8. Lecture Notes Series. 1993.

[124] K. T. Sturm. "Metric spaces of lower bounded curvature". *Exposition. Math.* 17.1 (1999), 35–47.

[125] В. А. Топоногов. «Римановы пространства кривизны, ограниченной снизу». *Успехи математических наук* 14.1 (85) (1959), 87—130.

[126] T. Toyoda. "An intrinsic characterization of five points in a CAT(0) space". *Anal. Geom. Metr. Spaces* 8.1 (2020), 114–165.

[127] C. Villani. *Optimal transport*. Vol. 338. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Old and new. 2009.

[128] A. Wald. „Begründung eiiner Koordinatenlosen Differentialgeometrie der Flächen". *Ergebnisse eines mathematischen Kolloquium* 6 (1935), 24–46.

[129] R. A. Wijsman. "Convergence of sequences of convex sets, cones and functions. II". *Trans. Amer. Math. Soc.* 123 (1966), 32–45.

[130] T. Yamaguchi. "Collapsing and pinching under a lower curvature bound". *Ann. of Math. (2)* 133.2 (1991), 317–357.