



universität
wien

MASTERARBEIT / MASTER'S THESIS

Titel der Masterarbeit / Title of the Master's Thesis

“Reshetnyak’s gluing theorem and its applications to billiards”

verfasst von / submitted by

Felix Rott, BSc

angestrebter akademischer Grad / in partial fulfilment of the requirements for the degree of

Master of Science (MSc)

Wien, 2020 / Vienna, 2020

Studienkennzahl lt. Studienblatt /
degree programme code as it appears on
the student record sheet:

A 066 821

Studienrichtung lt. Studienblatt /
degree programme as it appears on
the student record sheet:

Masterstudium Mathematik

Betreut von / Supervisor:

Univ.-Prof. Mag. Mag. Dr. Michael Kunzinger

Abstract

Alexandrov geometry can be regarded as a generalization of Riemannian geometry, in the sense that the differential structure is discarded completely, and the focus is shifted solely on the properties of the distance metric. The resulting spaces are called length spaces. As it turns out, while the assignment of an explicit curvature value is lost, one still retains the possibility of measuring curvature bounds. This is done by comparing geodesic triangles with appropriate triangles in the so called model spaces, i.e. the plane or a suitably curved sphere or hyperbolic space. An Alexandrov space is a length space that satisfies such a curvature bound, either from above or from below. This work only studies spaces with a curvature bound from above, called $\text{CAT}(k)$ spaces. The first chapter goes through elementary definitions and properties of length spaces, with a focus on gluing constructions. The second chapter introduces the model spaces, with which length spaces are compared to assign a curvature bound. In the third chapter, Reshetnyak's gluing theorem is proved, which ensures that when gluing two $\text{CAT}(k)$ spaces in a certain way, the resulting space satisfies the same curvature bound. The last chapter covers the collision theorem of semi-dispersing billiards as an application of the gluing theorem. This comparatively recent statement establishes an upper bound on the number of collisions of any trajectory that depends only on the number of walls and the minimal angle between them.

Abstrakt

Alexandrov-Geometrie kann als eine Verallgemeinerung der Riemannschen Geometrie betrachtet werden, in dem Sinne, dass die Differentialstruktur vollständig verworfen wird und der Schwerpunkt ausschließlich auf die Eigenschaften der Distanzmetrik verlagert wird. Die daraus resultierenden Räume werden als Längenräume bezeichnet. Wie sich herausstellt, geht zwar die Zuweisung eines expliziten Krümmungswertes verloren, aber man behält sich die Möglichkeit vor, Krümmungsschranken zu messen. Dies geschieht durch den Vergleich von geodätischen Dreiecken mit entsprechenden Dreiecken in den sogenannten Modellräumen, also der Ebene oder einer entsprechend gekrümmten Kugel oder einem hyperbolischen Raum. Ein Alexandrov-Raum ist ein Längenraum, der eine solche Krümmungsgrenze erfüllt, entweder von oben oder von unten. In dieser Arbeit werden nur Räume untersucht, deren Krümmung von oben begrenzt ist, so genannte $CAT(k)$ -Räume. Das erste Kapitel geht auf elementare Definitionen und Eigenschaften von Längenräumen ein, wobei der Schwerpunkt auf Verklebungskonstruktionen liegt. Das zweite Kapitel stellt die Modellräume vor, mit denen Längenräume verglichen werden, um eine Krümmungsschranke zuzuweisen. Im dritten Kapitel wird das Verklebungstheorem von Reshetnyak bewiesen, das sicherstellt, dass beim Verkleben von zwei $CAT(k)$ -Räumen auf eine bestimmte Art und Weise der resultierende Raum die gleiche Krümmungsschranke erfüllt. Das letzte Kapitel behandelt das Kollisionstheorem in semi-dispergierenden Billards als Anwendung des Verklebungstheorems. Diese vergleichsweise neue Aussage legt eine obere Schranke für die Anzahl der Kollisionen einer beliebigen Laufbahn fest, die nur von der Anzahl der Wände und dem minimalen Winkel zwischen diesen abhängt.

Contents

1	Introduction	1
1.1	Notation and conventions	1
1.2	Basic definitions	1
1.3	Constructing and characterizing length spaces	5
1.4	Gluing length spaces	6
2	A first look at curvature of length spaces	12
2.1	Introduction: curvature comparison in the plane	12
2.2	The model spaces	13
2.2.1	Euclidean space	13
2.2.2	The sphere	14
2.2.3	Hyperbolic space	16
2.2.4	The general model space M_k^n	18
2.3	Uniqueness of comparison triangles and Alexandrov's lemma	19
3	CAT(k) spaces	22
3.1	Basic properties	23
3.2	More comparisons and CAT(k) hierarchy	29
3.3	Final preparations and Reshetnyak's gluing theorem	31
4	Puff pastry and billiards	39
4.1	Preliminaries	39
4.2	Criteria for end-to-end convexity	41
4.3	The collision theorem	49

1 Introduction

In this chapter we give basic definitions and prerequisites. At first, we introduce length spaces and generalize the concepts of angles and triangles.

1.1 Notation and conventions

Here we briefly sum up the notation and conventions used throughout this work:

- (i) In a metric space (X, d) we denote by $B_r(x) := \{y \in X \mid d(x, y) < r\}$ and $\bar{B}_r(x) := \{y \in X \mid d(x, y) \leq r\}$ the open respectively closed ball around x with radius r . If $A \subseteq X$ and $x \in X$, then we set $d(x, A) := \inf\{d(x, a) \mid a \in A\}$. Similarly, if B is another subset of X , we set $d(A, B) := \inf\{d(a, b) \mid a \in A, b \in B\}$.
- (ii) Let $\sim$ be an equivalence relation on a set A . By $[a]_\sim := \{b \in A \mid a \sim b\}$ we denote the equivalence class of a . If it is clear from the context we will omit the $\sim$ and write $[a]$ instead. By $A/\sim := \{[a]_\sim \mid a \in A\}$ we denote the set of equivalence classes of A with respect to $\sim$. We will sometimes use $\tilde{A}$ instead of $A/\sim$ as well.
- (iii) A metric space X is called proper if every closed and bounded subset is compact. An equivalent statement is that for all $x \in X, r \geq 0$ the closed ball $\bar{B}_r(x)$ is compact.
- (iv) A map f between two metric spaces (X, d_X) and (Y, d_Y) is called nonincreasing (or short), if $d_X(x, y) \geq d_Y(f(x), f(y))$ holds for all $x, y \in X$.
- (v) If $s_1 = (x_1, \dots, x_n)$ and $s_2 = (y_1, \dots, y_m)$ denote two finite sequences, then we will call the sequence $(x_1, \dots, x_n, y_1, \dots, y_m)$ the concatenation of s_1 and s_2 , denoted by $s_1 \cdot s_2$.
- (vi) Let A be a subset of a topological space X . Then we write $\text{Int}(A)$ for its interior, $\text{Cl}(A)$ for its closure, and $\partial(A)$ for its boundary.

1.2 Basic definitions

In this section we mainly follow [BBI01, Ch. 2] and [BH99, Ch. I.1], as well as [Bal95, Ch. 1], [Hol13, Ch. 1], [Pap13, Ch. 1] and [Lan04].

1.2.1 Definition (Length structure). *A length structure is a triple (X, A, L) consisting of a topological space X , a subset A of all paths into X called (the class of) admissible paths and a map $L : A \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ called the length function satisfying the following properties:*

- (A1) *A is closed under restrictions, i.e. if $\gamma : [a, b] \rightarrow X$ is admissible, then any restriction $\gamma|_{[c, d]} : [c, d] \rightarrow X$ is again admissible ($a \leq c < d \leq b$).*
- (A2) *A is closed under concatenation, i.e. if $\gamma_1 : [a, b] \rightarrow X$ and $\gamma_2 : [c, d] \rightarrow X$ are two admissible paths such that $\gamma_1(b) = \gamma_2(c)$, then also their concatenation $\gamma_1 \cdot \gamma_2 : [a, b + d - c] \rightarrow X$ defined by*

$$\gamma_1 \cdot \gamma_2(t) := \begin{cases} \gamma_1(t), & \text{if } t \in [a, b] \\ \gamma_2(t - b + c), & \text{if } t \in [b, b + d - c] \end{cases} \quad (1.2.1)$$

is admissible.

- (A3) *A is closed under linear reparameterizations, i.e. if $\gamma : [a, b] \rightarrow X$ is admissible and $\varphi : [c, d] \rightarrow [a, b]$, $\varphi(t) = \frac{a-b}{c-d}(t-d) + b$ or $\varphi(t) = \frac{b-a}{c-d}(t-d) + a$ then $\gamma \circ \varphi : [c, d] \rightarrow X$ is again admissible.*

The length function has to, in some sense, be compatible with these three properties and additionally respect the original topology of X in a certain way. More concretely:

- (L1) The length is additive, i.e. if $\gamma : [a, b] \rightarrow X$ is admissible, then $L(\gamma) = L(\gamma|_{[a, c]}) + L(\gamma|_{[c, b]})$ for any $c \in [a, b]$.
- (L2) The length is invariant under linear reparameterizations, i.e. if $\gamma : [a, b] \rightarrow X$ is admissible and φ as in (A3), then $L(\gamma) = L(\gamma \circ \varphi)$.
- (L3) The length depends continuously on the parameter of the path, i.e. if $\gamma : [a, b] \rightarrow X$ is admissible and $L(\gamma) < \infty$, then $t \mapsto L(\gamma|_{[a, t]})$, $t \in [a, b]$ is continuous. In addition, we assume every constant path to be of length 0.
- (L4) The length has to respect the topology of X in the following way: for any $x \in X$ and any neighbourhood U_x of x the length of a path connecting x with a point outside of U_x has to be positive, i.e.

$$\inf\{L(\gamma) \mid A \ni \gamma : [a, b] \rightarrow X, \gamma(a) = x, \gamma(b) \in X \setminus U_x\} > 0. \quad (1.2.2)$$

In order to ensure that we can induce a metric structure on X , we will assume X to be Hausdorff.

1.2.2 Definition (Length metric and length space). *Let (X, A, L) be a length structure. This length structure induces a length metric d_L on X by setting the distance between two points as the infimum over all admissible paths connecting them, i.e.*

$$d_L(x, y) := \inf\{L(\gamma) \mid A \ni \gamma : [a, b] \rightarrow X, \gamma(a) = x, \gamma(b) = y\}. \quad (1.2.3)$$

A metric space (X, d) is called a length space if there is a length structure (X, A, L) such that its induced length metric is d .

If there is no risk of confusion, we will follow the usual convention of writing X instead of (X, d) . Also, we will usually denote all metrics with d when it is clear from the context to which space they belong. E.g. when working with model spaces, all their points will be denoted with a bar to remove ambiguity.

1.2.3 Definition (Shortest paths and (local) geodesics). *Let X be a length space.*

- (i) A shortest path between x and y in X is an admissible path $\gamma : [a, b] \rightarrow X$ such that $L(\gamma) = d(x, y)$.
- (ii) A geodesic is a shortest path that is parameterized by arclength, i.e. an admissible path $\gamma : [0, l] \rightarrow X$ such that $d(\gamma(s), \gamma(t)) = |s - t|$ for all $s, t \in [0, l]$. In other words, a geodesic is an isometry from a compact connected subset of $\mathbb{R}$ onto its image.
- (iii) An admissible path $\gamma : [0, l] \rightarrow X$ is called a local geodesic if it is locally a shortest path that is parameterized by arclength, i.e. for all $t \in [0, l]$ there exists a subinterval $I \subseteq [a, b]$ containing t such that $\gamma|_I : I \rightarrow X$ is a geodesic.

Note that because of (A3) and (L1), the restriction of a shortest path respectively a (local) geodesic is again a shortest path respectively a (local) geodesic. We also observe that geodesics from Riemannian geometry are in general only local geodesics with respect to our definition: e.g. on the sphere, any arc of a great circle is a geodesic in the Riemannian sense, but it is a geodesic as in 1.2.3 only if its length is $< \pi$. Similar to real intervals, we will denote the image

$\{\gamma(t) \mid t \in [0, l]\}$ of a geodesic $\gamma : [0, l] \rightarrow X$, where $\gamma(0) = x$ and $\gamma(l) = y$ as $[x, y]$ and we will call this image a geodesic segment joining x and y . Note that there may be many geodesics joining two points, so we will only use this notation if the choice of a geodesic does not matter or if there is no choice at all. If $[x, y]$ and $[y, z]$ are two geodesic segments, then analogously to the concatenation of paths we will denote the concatenation of the segments by $[x, y] \cdot [y, z]$. However, this only makes sense if the result is actually a geodesic segment again.

As in Riemannian geometry, geodesics are an extremely important subclass of paths that do enjoy a wide range of nice properties. Also, many of the constructions below only work when considering geodesics instead of arbitrary (admissible) paths. We also remark here that it is important to require the infimum in the definition of d_L instead of the minimum. Consider e.g. the Euclidean plane with the origin removed. This is clearly a length space but for points opposite of each other with respect to the origin there is no shortest path connecting them. The property that such a shortest path always exists is worth a definition.

1.2.4 Definition (Geodesic and r -geodesic length space). *Let (X, d) be a length space.*

- (i) *(X, d) is called a complete length space (or a geodesic space or geodesically complete) if for all $x, y \in X$ there is a path of length $d(x, y)$ connecting them. In other words, the infimum is always achieved by an element of A .*
- (ii) *Let $r > 0$. (X, d) is called r -geodesic if for all points $x, y \in X$ with $d(x, y) < r$ there exists a geodesic between them.*

1.2.5 Remark (Geodesic completeness versus metric completeness). There is a generalization of the well known Hopf-Rinow theorem from Riemannian geometry to length spaces. It states (among other things) that a locally compact length space is geodesically complete if and only if it is complete as a metric space, i.e. every Cauchy sequence converges. Since the length spaces we are working with in general need not be locally compact we will usually call a complete length space in the above sense a geodesic space, or say that X is geodesic.

1.2.6 Definition (Convex subsets). *Let X be a length space.*

- (i) *A subset $A \subseteq X$ is said to be convex if for all $x, y \in A$ there is a geodesic joining them and all such geodesic segments are entirely contained in A .*
- (ii) *Similar to r -geodesicity, we have that a subset $A \subseteq X$ is r -convex, if for all $x, y \in A$ with $d(x, y) < r$ there is a geodesic joining them and all such geodesic segments are entirely contained in A .*

Before we can assign a curvature bound to our length spaces we first have to introduce the concepts of angles and triangles into the setting of length spaces.

1.2.7 Definition (Triangles and comparison triangles). *Let X be a length space.*

- (i) *A (geodesic) triangle in X is a collection of three (distinct) points x, y, z together with three geodesic segments $[x, y], [y, z], [z, x]$. We will denote such a triangle by $\Delta([x, y], [y, z], [z, x])$ or by $\Delta(x, y, z)$ if the choice of geodesic segments does not matter or geodesics in X are unique.*
- (ii) *A comparison triangle for $\Delta([x, y], [y, z], [z, x])$ is a triangle in the Euclidean plane with corresponding side lengths, i.e. a collection of three points $\bar{x}, \bar{y}, \bar{z} \in \mathbb{R}^2$ such that $d_{\mathbb{R}^2}(\bar{x}, \bar{y}) = d(x, y)$, $d_{\mathbb{R}^2}(\bar{y}, \bar{z}) = d(y, z)$ and $d_{\mathbb{R}^2}(\bar{z}, \bar{x}) = d(z, x)$. We will write $\bar{\Delta}(\bar{x}, \bar{y}, \bar{z})$ for the comparison triangle of $\Delta([x, y], [y, z], [z, x])$ and $[\bar{x}, \bar{y}]$ for the straight line segment connecting $\bar{x}$ and $\bar{y}$. We will show later that the comparison triangle is unique up to isometry.*

- (iii) If a point $p \in X$ is an element of the union of $[x, y], [y, z], [z, x]$, i.e. “lies on a side of the triangle”, then we write $p \in \Delta([x, y], [y, z], [z, x])$. We say a triangle $\Delta([x, y], [y, z], [z, x])$ lies in a subset $U \subseteq X$ if $x, y, z \in U$ and $[x, y], [y, z], [z, x] \subseteq U$ and we will write (somewhat inaccurately) $\Delta([x, y], [y, z], [z, x]) \in U$.

1.2.8 Definition (Alexandrov angle and comparison angle). *Let (X, d) be a length space.*

- (i) Let $\bar{\Delta}(\bar{x}, \bar{y}, \bar{z})$ be a comparison triangle for a triangle $\Delta([x, y], [y, z], [z, x])$ in X . The comparison angle at x between y and z is the interior angle of the comparison triangle $\bar{\Delta}(\bar{x}, \bar{y}, \bar{z})$ at $\bar{x}$ and is denoted by $\bar{\angle}_x(y, z)$. The comparison angle can also be defined explicitly:

$$\bar{\angle}_x(y, z) := \arccos \left(\frac{d(x, y)^2 + d(x, z)^2 - d(y, z)^2}{2d(x, y)d(x, z)} \right). \quad (1.2.4)$$

- (ii) Let $\gamma : [0, l] \rightarrow X$ and $\gamma' : [0, l'] \rightarrow X$ be two geodesics with $p := \gamma(0) = \gamma'(0)$. Then the Alexandrov angle between γ and γ' at p is defined as

$$\angle_p(\gamma, \gamma') := \limsup_{t, t' \rightarrow 0} \bar{\angle}_{\gamma(0)}(\gamma(t), \gamma'(t')). \quad (1.2.5)$$

We will often write angle instead of Alexandrov angle. If $[p, x]$ and $[p, y]$ are the geodesic segments corresponding to γ and γ' , respectively, we will also denote the Alexandrov angle by $\angle_p([p, x], [p, y])$, if (more) suitable for our situation.

1.2.9 Remark (Remark on angles). (i) One motivation for the Alexandrov angle comes from Riemannian geometry. On a Riemannian manifold (M, g) the angle between two smooth curves $c_1, c_2 : \mathbb{R} \rightarrow M$ intersecting at $p := c_1(0) = c_2(0)$ is defined as

$$\arccos \left(\frac{g_p(c'_1(0), c'_2(0))}{\sqrt{g_p(c'_1(0), c'_1(0))g_p(c'_2(0), c'_2(0))}} \right), \quad (1.2.6)$$

i.e. the “usual” angle between their tangent vectors at p . Looking at smaller and smaller triangles in the definition of the Alexandrov angle via the limit this construction resembles the definition of a tangent vector.

- (ii) Using the explicit description of the comparison angle one can obtain an alternative definition of the Alexandrov angle. Setting $p := \gamma(0) = \gamma'(0)$, we have $d(p, \gamma(t)) = t$ and $d(p, \gamma'(s)) = s$. Thus,

$$\angle_p(\gamma, \gamma') = \limsup_{t, s \rightarrow 0} \arccos \left(\frac{t^2 + s^2 - d(\gamma(t), \gamma'(s))^2}{2ts} \right). \quad (1.2.7)$$

This also shows that the angle varies between 0 and π (if it exists).

- (iii) The angle remains unchanged when considering initial segments of the geodesics. More precisely, if $[p, x]$ and $[p, y]$ are two distinct geodesic segments, then we have $\angle_p([p, x], [p, y]) = \angle_p([p, q_1], [p, q_2])$ for all $q_1 \in [p, x], q_2 \in [p, y]$.
- (iv) If p lies in the interior of a geodesic segment $[x, y]$ then the angle at p between the restriction of the two segments is π . More precisely, if $p \in [x, y]$ and $[x, y] = [x, p] \cdot [p, y]$, we have $\angle_p([x, p], [p, y]) = \pi$.

- (v) Last but definitely not least we want to mention that the angle also satisfies the triangle inequality. More precisely, let $[p, x], [p, y], [p, z]$ be three geodesic segments issuing from the same point $p \in X$. Then

$$\angle_p([p, x], [p, z]) \leq \angle_p([p, x], [p, y]) + \angle_p([p, y], [p, z]). \quad (1.2.8)$$

For a proof of this statement, see e.g. [BH99, Ch. I, Prop. 1.14].

1.3 Constructing and characterizing length spaces

In this section, we briefly want to discuss how one can obtain a length space from an arbitrary metric space. In other words, any metric induces a length function on the class of all paths. This gives a length structure which in turn then induces a length metric. First, we need a suitable notion of the length of a path in an arbitrary metric space. The following definition resembles the (very intuitive!) construction from analysis where one approximates the length of a real path by straight line segments. Again, we follow mainly [BBI01, Ch. 2], [Bal95, Ch. 1], [Hol13, Ch. 1], [Pap13, Ch. 1 & 2] and [Lan04]. We will only give references for the proofs in this section.

1.3.1 Definition (Total variation and variational length). *Let (X, d) be a metric space, $\gamma : [a, b] \rightarrow X$ a path and $\sigma = \{t_0, \dots, t_n\} \subseteq [a, b]$ such that $t_0 = a, t_n = b$ and $t_i < t_j \forall i < j$. We will call such a subset a partition of $[a, b]$.*

- (i) *The total variation of γ with respect to the partition σ is defined as*

$$V_\sigma(\gamma) := \sum_{i=0}^{n-1} d(\gamma(t_i), \gamma(t_{i+1})). \quad (1.3.1)$$

- (ii) *The variational length of γ is the supremum of all total variations of γ , i.e.*

$$L_d(\gamma) := \sup\{V_\sigma \mid \sigma = \{t_0, \dots, t_n\} \text{ a partition of } [a, b], n \in \mathbb{N}\}. \quad (1.3.2)$$

Clearly, by applying the triangle inequality multiple times, the variational length satisfies

$$L_d(\gamma) \geq d(\gamma(a), \gamma(b)). \quad (1.3.3)$$

For a proof of the following statement, see e.g. [BBI01, Prop. 2.3.4].

1.3.2 Proposition (Variational length structure). *Let (X, d) be a metric space and C the class of all (continuous) paths in X . Then (X, C, L_d) is a length structure. In particular, (X, d_{L_d}) is a length space.* $\square$

The following characterization of length spaces will be useful. It is a corollary of [BBI01, Props. 2.3.12 & 2.4.1].

- 1.3.3 Proposition** (Characterizing length spaces). *(i) (X, d) is a length space if and only if for all $\varepsilon > 0$ and all $x, y \in X$ there is a path γ joining x and y with $L_d(\gamma) < d(x, y) + \varepsilon$.*
(ii) (X, d) is a geodesic space if and only if for all $x, y \in X$ there is a path γ joining x and y with $L_d(\gamma) = d(x, y)$. $\square$

Another convenient characterization is via the existence of midpoints.

1.3.4 Definition (Midpoints and ε -midpoints). *Let (X, d) be a metric space.*

- (i) *A point $y \in X$ is called a midpoint for $x, z \in X$ if $d(x, y) = d(y, z) = \frac{1}{2}d(x, z)$.*
- (ii) *A point $y \in X$ is called ε -midpoint for $x, z \in X$ if $|2d(x, y) - d(x, z)| \leq \varepsilon$ and $|2d(y, z) - d(x, z)| \leq \varepsilon$.*

Note that ε -midpoints are often called approximate midpoints, especially if the value of ε is not important. For a proof of the following proposition, see e.g. [BBI01, Lemmata 2.4.8, 2.4.10 & Thm. 2.4.16].

1.3.5 Proposition (Characterizing length spaces via midpoints). *Let (X, d) be a complete metric space.*

- (i) *X is a geodesic space if and only if for all pairs of points there exists a midpoint.*
- (ii) *X is a length space if and only if for all pairs of points and all $\varepsilon > 0$ there is an ε -midpoint.*

□

1.3.6 Remark (Completeness is not restrictive). Note that the completeness of X is only required for the “only if” part of the above proposition, i.e. if X is a not necessarily complete geodesic space, then still every pair of points has a midpoint. This is quite fortunate since we will mainly make use of this direction.

1.4 Gluing length spaces

In this section we consider general gluing constructions for length spaces. It is a very natural question to ask what properties are preserved under gluing. We will discuss among other things the disjoint union of length spaces and the quotient pseudometric. We will show that a (certain) quotient of a length space is again a length space and look at the gluing of triangles and gluing along isometric subspaces. Here, the contents will be taken from [BH99, Ch. I.5 & II.4] and [BBI01, Ch. 3] as well as [Stu06, Ch. 2 & 3], [Pau04, Ch. 2], [Hol13, Ch. 1] and [Lan04].

1.4.1 Definition (Disjoint union and corresponding metric). *Let I be an index set.*

- (i) *Formally, the disjoint union of a family of sets $(A_i)_{i \in I}$ is defined as $A := \bigsqcup_{i \in I} A_i := \{(a, i) \mid a \in A_i, i \in I\}$. We will usually consider sets that are already disjoint and so we omit the second component of the elements of A .*
- (ii) *Let $(X_i, d_i)_{i \in I}$ be a family of length spaces. We define a metric d on their disjoint union $X := \bigsqcup_{i \in I} X_i$ by setting*

$$d(x, y) := \begin{cases} d_i(x, y), & \text{if } x, y \in X_i, \\ \infty, & \text{otherwise.} \end{cases} \quad (1.4.1)$$

While this construction is elementary and can always be done, it does not say much about the respective spaces and how they interact with each other. In general topology there is the notion of gluing along an equivalence relation. We aim to develop a similar concept that is compatible with the original metric of the space. This construction is known as the quotient semi-metric.

1.4.2 Definition (Quotient semi-metric). *Let (X, d) be a metric space, $\sim$ an equivalence relation on X and set $\tilde{X} := X/\sim$. Then the map $\tilde{d} : \tilde{X} \times \tilde{X} \rightarrow \mathbb{R}$ defined by*

$$\tilde{d}([x], [y]) := \inf \left\{ \sum_{i=1}^n d(x_i, y_i) \mid x \sim x_1, y \sim y_n, x_{i+1} \sim y_i \ \forall i = 1, \dots, n-1, n \in \mathbb{N} \right\} \quad (1.4.2)$$

is called the quotient semi-metric (associated to d).

Especially in the setting of (geodesic) length spaces this definition is rather intuitive: the distance of x and y is approximated by pieces of paths glued along their endpoints. We denote by $S_{x,y}$ the set of all sequences over which we take the infimum in (1.4.2), i.e. $S_{x,y} := \{(x_1, y_1, x_2, \dots, y_n) \mid x \sim x_1, y \sim y_n, x_{i+1} \sim y_i \ \forall i = 1, \dots, n-1, n \in \mathbb{N}\}$. The elements $s \in S_{x,y}$ are also called n -chains and the length of a specific n -chain will be denoted by $l(s)$, i.e. if $s = (x_1, y_1, x_2, \dots, y_n)$, then $l(s) := \sum_{i=1}^n d(x_i, y_i)$. The quotient semi-metric then takes the abbreviated form $\tilde{d}([x], [y]) = \inf_{s \in S_{x,y}} l(s)$. Note that since $(x, y) \in S_{x,y}$, we always have $\tilde{d}([x], [y]) \leq d(x, y)$. In other words, if $\pi : X \rightarrow \tilde{X}$ denotes the natural projection then $\tilde{d} \circ (\pi \times \pi) \leq d$. In particular, if $\tilde{d}$ is a metric then $L_{\tilde{d}}(\pi \circ \gamma) \leq L_d(\gamma)$ for any (admissible) path $\gamma : [a, b] \rightarrow X$ since

$$V_\sigma(\pi \circ \gamma) = \sum_{i=0}^{n-1} \tilde{d}([\gamma(t_i)], [\gamma(t_{i+1})]) \leq \sum_{i=0}^{n-1} d(\gamma(t_i), \gamma(t_{i+1})) = V_\sigma(\gamma) \quad (1.4.3)$$

for any partition σ of $[a, b]$.

1.4.3 Proposition (Quotient semi metric is well named). *Let (X, d) be a metric space. Then $(\tilde{X}, \tilde{d})$ is a semi-metric space, i.e. $\tilde{d}$ satisfies all properties of a metric except that different points have positive distance from each other.*

Proof. Clearly, $\tilde{d}$ is symmetric (since d is) and $\tilde{d}([x], [x]) = 0$. Let $\varepsilon > 0$ and let $s_1 = (x_1, y_1, x_2, \dots, y_k)$ and $s_2 = (y_k, z_k, y_{k+1}, \dots, z_n)$ be two sequences in $S_{x,y}$ and $S_{y,z}$ such that $l(s_1) < \tilde{d}([x], [y]) + \frac{\varepsilon}{2}$ and $l(s_2) < \tilde{d}([y], [z]) + \frac{\varepsilon}{2}$, respectively. Then clearly

$$l(s_1 \cdot s_2) = l(s_1) + l(s_2) < \tilde{d}([x], [y]) + \tilde{d}([y], [z]) + \varepsilon. \quad (1.4.4)$$

Since $s_1 \cdot s_2 \in S_{x,z}$, we have

$$l(s_1 \cdot s_2) \geq \inf_{s \in S_{x,z}} l(s) = \tilde{d}([x], [z]). \quad (1.4.5)$$

Since ε was arbitrary, $\tilde{d}([x], [z]) \leq \tilde{d}([x], [y]) + \tilde{d}([y], [z])$ follows.

Finally, observe that $\tilde{d}$ may fail to be an actual metric. Intuitively, one can observe that if two points in X can be approximated to arbitrary precision by equivalent points, then their distance will be zero. Specifically, consider the real line (with the usual Euclidean metric) with all rational points identified. Then the quotient semi-metric is the constant zero function. $\square$

1.4.4 Remark (Turning a semi-metric into a metric). There is an obvious way how to turn any semi-metric space into a metric space: if (X, d) is a semi-metric space, define an equivalence relation R via $xRy \Leftrightarrow d(x, y) = 0$. Reflexivity and symmetry for R are directly inherited from the semi-metric. And if $d(x, y) = d(y, z) = 0$ it follows that $0 \leq d(x, z) \leq d(x, y) + d(y, z) = 0$. Thus, R is an equivalence relation on X and we can form the quotient X/R . This can be made into a metric space: simply define $d_R([x], [y]) := d(x, y)$. This is well defined: Let $x_1, x_2, y_1, y_2 \in X$ be two representatives from the equivalence classes $[x]$ and $[y]$, respectively. Then

$$d_R([x_1], [y_1]) = d(x_1, y_1) \leq d(x_1, x_2) + d(x_2, y_1) = d(x_2, y_1).$$

But also

$$d_R([x_2], [y_1]) = d(x_2, y_1) \leq d(x_2, x_1) + d(x_1, y_1) = d(x_1, y_1).$$

Thus, $d(x_1, y_1) = d(x_2, y_1)$ whenever x_1, x_2 lie in the same equivalence class. Symmetry of the metric gives the same for $[y]$. So $(X/R, d_R)$ is a metric space which we will denote by $X/\bar{d}$. In particular, let (X, d) be a metric space and $(\tilde{X}, \tilde{d})$ the corresponding semi-metric space. Then we can identify points of vanishing $\tilde{d}$ distance in $\tilde{X}$ to obtain a metric space again. This space will be denoted by $X/\bar{d}$.

1.4.5 Remark ($X/\bar{d}$ may differ wildly from $X/\sim$). It may happen that there are fewer equivalence classes in $X/\bar{d}$ than in $X/\sim$. In other words, the two spaces need not be homeomorphic. Consider the counterexample in 1.4.3: $\mathbb{R}/\mathbb{Q}$ is not a T_1 -space: if we denote by $\pi : \mathbb{R} \rightarrow \mathbb{R}/\mathbb{Q}$ the natural projection, then for any $x \in \mathbb{Q}$ we have $\pi^{-1}[\{\pi(x)\}] = \mathbb{Q}$, which is not closed in $\mathbb{R}$, hence the singleton $\{\pi(x)\}$ is not closed in $\mathbb{R}/\mathbb{Q}$. $(\mathbb{R}/\mathbb{Q})/\bar{d}$, however, is a (one point, admittedly,) metric space.

As a first step we show that the quotient semi-metric yields especially interesting results in the case of length spaces.

1.4.6 Proposition (Quotient semi-metric in length spaces is well behaved). *Let (X, d) be a length space, $(\tilde{X}, \tilde{d})$ the associated quotient semi-metric space. If $(\tilde{X}, \tilde{d})$ is a metric space then it is a length space.*

Proof. Let $\varepsilon > 0$ and let $(x_1, y_1, x_2, \dots, y_n)$ be an n -chain such that $\sum_{i=1}^n d(x_i, y_i) < \tilde{d}([x], [y]) + \frac{\varepsilon}{2}$. Since (X, d) is a length space, for all x_i, y_i there is a path γ connecting them with $L_d(\gamma_i) < d(x_i, y_i) + \frac{\varepsilon}{2n}$. Set $\tilde{\gamma}_i := \pi \circ \gamma$ and $\tilde{\gamma} := \tilde{\gamma}_1 \cdot \dots \cdot \tilde{\gamma}_n$. Then we have

$$L_{\tilde{d}}(\tilde{\gamma}_1 \cdot \dots \cdot \tilde{\gamma}_n) \leq L_d(\gamma_1 \cdot \dots \cdot \gamma_n) \leq \sum_{i=1}^n \left(d(x_i, y_i) + \frac{\varepsilon}{2n} \right) < \tilde{d}([x], [y]) + \varepsilon. \quad (1.4.6)$$

In summary, for any $\varepsilon > 0$ we found a path $\tilde{\gamma}$ connecting $[x]$ and $[y]$ such that $L_{\tilde{d}}(\tilde{\gamma}) < \tilde{d}([x], [y]) + \varepsilon$. Thus, by 1.3.3 $(\tilde{X}, \tilde{d})$ is a length space. $\square$

1.4.7 Remark (Geodesic quotient may fail to be geodesic). Intuitively, one may already think that if (X, d) is a geodesic space then $(\tilde{X}, \tilde{d})$ need not be a geodesic space: the definition of the quotient semi-metric adds another infimum to the metric which could prevent us from always finding a shortest path. Indeed, if between two fixed points we have infinitely many geodesics of fixed distance from each other and then identify bigger and bigger neighbourhoods of the midpoints of these geodesics, there may not be a shortest path between the original endpoints. Inspired by [Pet19], we give the following concrete counterexample: consider the Hilbert space of square-summable real sequences $l^2(\mathbb{N}) := \{x = (x_1, x_2, \dots) \mid x_i \in \mathbb{R}, \sum_{i=1}^{\infty} |x_i|^2 < \infty\}$. Let S be the unit sphere in $l^2(\mathbb{N})$, i.e. $S := \{x \in l^2(\mathbb{N}) \mid \sum_{i=1}^{\infty} |x_i|^2 = 1\}$. This is one of the more natural ways to imagine an infinite dimensional sphere. Using the metric in (2.2.2) below, we have that (S, d) is a geodesic space. We denote by e_i the sequence that has 1 as i -th component and zero otherwise, i.e. $e_i := (\delta_{i,1}, \delta_{i,2}, \dots)$. Then clearly $e_i \in S$ for all $i \in \mathbb{N}$. Using the constructions from 2.2.2, the path $c_i : [0, \pi] \rightarrow S$ defined by $c_i(t) := \cos(t)e_1 + \sin(t)e_i$ is a geodesic from e_1 to $-e_1$ for all $i \geq 2$ with e_i being the midpoint of c_i . Clearly, $d(e_i, e_j) = \frac{\pi}{2}$ for all $i \neq j$ and so we have a collection of infinitely many geodesics whose midpoints have a fixed distance from each

other. Now set $r_i := 1 - \frac{1}{i}$ and consider the equivalence relation on S that identifies all points in $B_{\frac{r_i}{2}}(e_i)$ for all $i \geq 2$. Note that the resulting equivalence classes are disjoint in the sense that no $x \in S$ belongs to two different balls because the radii are small enough. Then we clearly have $\tilde{d}(e_1, -e_1) = \pi - 1$, but no path achieves this distance. Thus, $(\tilde{S}, \tilde{d})$ is not a geodesic space.

We can further develop the above constructions. Instead of identifying points in a single metric space, we can glue together two (or several, even infinitely many) metric spaces. The most intuitive way is to glue them along isometric subsets.

1.4.8 Definition (Amalgamation of metric spaces). *Let $(X_i, d_i)_{i \in I}$ be a family of metric spaces and $(A_i, d_i|_{A_i \times A_i})_{i \in I}$ a family of closed isometric subspaces. Let A be a metric space isometric to each A_i via the isometry $\phi_i : A \rightarrow A_i$. Equip the disjoint union $\bigsqcup_{i \in I} X_i$ with the metric from 1.4.1. Let $\sim$ be the equivalence relation generated by the condition $\phi_i(a) \sim \phi_j(a)$ for all $i, j \in I, a \in A$. More precisely, define $\sim$ via*

$$x \sim y :\Leftrightarrow x = y \text{ or } \exists a \in A : \phi_i(a) = x, \phi_j(a) = y. \quad (1.4.7)$$

Then the quotient of $\bigsqcup_{i \in I} X_i$ with respect to $\sim$ is called the amalgamation of the family X_i with respect to A and is denoted by $\bigsqcup_A X_i$.

1.4.9 Proposition (Properties of the amalgamation). *Let $X := \bigsqcup_A X_i$ be the amalgamation of a family of metric spaces $(X_i, d_i)_{i \in I}$.*

(i) *The quotient semi-metric with respect to the amalgamation relation on X takes the form*

$$\tilde{d}([x], [y]) = \begin{cases} d_i(x, y), & \text{if } x, y \in X_i, \\ \inf_{a \in A} \{d_i(x, \phi_i(a)) + d_j(\phi_j(a), y)\}, & \text{if } x \in X_i, y \in X_j. \end{cases} \quad (1.4.8)$$

In particular, $(X, \tilde{d})$ is always a metric space.

(ii) *If I is finite and X_i is proper for all $i \in I$, then X is proper.*

(iii) *If X_i is a geodesic space for all $i \in I$ and A is proper, then X is a geodesic space.*

(iv) *If X_i is a length space for all $i \in I$, then X is a length space. Furthermore, $\tilde{d}$ is the unique length metric on X such that its restriction to each X_i is d_i .*

Proof. (i) Denote by d the metric of the disjoint union $\bigsqcup_{i \in I} X_i$. Let $x, y \in X_m$. Then $\tilde{d}([x], [y]) = \inf_{s \in S_{x,y}} l(s)$ by definition of the quotient semi-metric. If $d_m(x, y) = \infty$, then by the triangle inequality we infer that for any other $z \in X_m$ (at least one of) $d_m(x, z) = \infty$ or $d_m(y, z) = \infty$ hold. Together with the fact that the equivalence relation does not identify points within the same space, we conclude that $\tilde{d}([x], [y]) = \infty$ as well. So now assume $d_m(x, y) < \infty$ and let $s = (x_1, y_1, x_2, \dots, y_n) \in S_{x,y}$ with $l(s) < \infty$. This exists since $(x, y) \in S_{x,y}$ and so $\tilde{d}([x], [y]) \leq d(x, y) < \infty$. We now argue that if s leaves X_m then there is a shorter sequence: Let i be the first index where the n -chain s leaves X_m , i.e. $x_i, y_i \in X_m$ and say $x_{i+1}, y_{i+1} \in X_j$. Clearly, points with the same index have to belong to the same X_i , otherwise d would return infinite distance by definition of the disjoint union metric (1.4.1). We have $y_i \sim x_{i+1}$ and since clearly $y_i \neq x_{i+1}$ (they belong to different spaces) there exists an $a \in A$ such that $\phi_m(a) = y_i$ and $\phi_j(a) = x_{i+1}$. Now suppose $y_{i+1} \in X_j \setminus A_j$. s ultimately has to return to X_m and can leave X_j only through A_j . Let $i' > i + 1$ be the first index where s leaves X_j , i.e. $x_{i'}, y_{i'} \in X_j$ and $x_{i'+1}, y_{i'+1}$ lie in another space, say $X_{j'}$. Since $\{x_{i+1}, y_{i+1}, \dots, x_{i'}, y_{i'}\} \subseteq X_j$ the only possibility for relatedness between them is equality. Thus, the triangle inequality gives $\sum_{k=i+1}^{i'} d(x_k, y_k) =$

$\sum_{k=i+1}^{i'} d_j(x_k, y_k) \geq d_j(x_{i+1}, y_{i'})$ and so we can omit the “detour” through $X_j \setminus A_j$ to obtain a lengthwise shorter sequence than s . Repeating this for all other detours, we have established that either s stays completely in X_m or it “jumps around” between different A_i ’s. But since the A_i ’s are isometric, by the triangle inequality again we obtain a (possibly) shorter sequence fully staying in X_m . More precisely, consider $x_k \in X_m, y_k \in A_m \subseteq X_m, x_{k+1}, y_{k+1} \in A_j, y_k \sim x_{k+1}$. Using the triangle inequality and the fact that ϕ_j is an isometry, we compute

$$\begin{aligned} d(x_k, y_k) + d(x_{k+1}, y_{k+1}) &= d_m(x_k, \phi_m(a)) + d_j(\phi_j(a), y_{k+1}) \\ &= d_m(x_k, \phi_m(a)) + d_A(a, \phi_j^{-1}(y_{k+1})) \\ &= d_m(x_k, \phi_m(a)) + d_m(\phi_m(a), \phi_m(\phi_j^{-1}(y_{k+1}))) \\ &\geq d_m(x_k, \phi_m(\phi_j^{-1}(y_{k+1}))). \end{aligned}$$

We can repeat this process to obtain $l(s) \geq d(x, y)$ in the end. Since s was arbitrary, the first case in (1.4.8) follows.

Now let $x \in X_m$ and $y \in X_n$ and let $s \in S_{x,y}$. By the above consideration, we only have to consider chains that start in X_m , jump once to A_n via A_m and then stay in X_n . Using the triangle inequality we shorten the first part of the sum to the distance of x and the first point where the sequence leaves A_m , say after y_i . Similarly, we shorten the final part of the sequence to the distance between the point where s first enters A_n , say x_j ($i < j$), and y . Hence, the length of our improved sequence is given by

$$d(x, y_i) + \sum_{k=i+1}^{j-1} d(x_k, y_k) + d(x_j, y), \quad (1.4.9)$$

where x_k, y_k belong to the same X_l . By applying the triangle inequality and various isometries and their inverse functions, we can shorten this even further to $d(x, y_i) + d(x_j, y)$, where $y_i = \phi_m(a)$ and $x_j = \phi_n(a)$ for some $a \in A$. Summing up, any sequence s can be “improved” by this construction, so we only need to take our infimum over expressions of this form, giving (1.4.8).

To conclude the proof of (i), we have to show that $\tilde{d}$ actually is a metric on X , i.e. $\tilde{d}([x], [y]) > 0$ for $[x] \neq [y]$. If $x \in X_n \setminus A_n$ and consequently $[x] = \{x\}$ then $\tilde{d}([x], [y]) > 0$ follows since A_n is closed. Finally, let $x \in A_m, y \in A_n$ and $[x] \neq [y]$. Since both A_m and A_j are isometric to A , $d([x], [y])$ can be identified with a distance in A , which has to be positive because A is a metric space.

(ii) By construction, we have $\tilde{d}|_{X_i \times X_i} = d_i$. Thus, the intersection of X_i with a closed and bounded subset $B \subseteq X$ is closed and bounded in X_i . Now if (X_i, d_i) is a proper metric space for all $i \in I$ and I is finite, then any closed and bounded subset of X is a finite union of compact sets and hence itself compact, implying that X is a proper metric space.

(iii) Clearly, if $x, y \in X_i$, then $\tilde{d}([x], [y]) = d_i(x, y)$ and one can take the (projection of the) original geodesic between them as a shortest path. Also, if say $x \in X_i$ and $y \in A_j$, then A_j can be identified with A_i and we find a shortest path between them since X_i is a geodesic space. Let $[x] \neq [y] \in X$ and say $x \in X_i \setminus A_i, y \in X_j \setminus A_j$. First observe that by definition of $\tilde{d}([x], [y])$, there is an $a \in A$ such that $d_i(x, \phi_i(a)) \leq \tilde{d}([x], [y])$: suppose to the contrary this is not true, i.e. for all $a \in A$ we have $d_i(x, \phi_i(a)) > \tilde{d}([x], [y])$. Then we compute

$$\begin{aligned}\tilde{d}([x], [y]) &= \inf_{a \in A} \{d_i(x, \phi_i(a)) + d_j(\phi_j(a), y)\} \geq \inf_{a \in A} \{\tilde{d}([x], [y]) + d_j(\phi_j(a), y)\} \\ &= \tilde{d}([x], [y]) + \inf_{a \in A} \{d_j(\phi_j(a), y)\} > \tilde{d}([x], [y]),\end{aligned}$$

since $\inf_{a \in A} \{d_j(\phi_j(a), y)\} = \inf_{b \in A_j} \{d_j(b, y)\} > 0$ if $y \in X_j \setminus A_j$ since A_j is closed. The above considerations imply that the set $D_x := \{a \in A \mid d_i(x, \phi_i(a)) \leq \tilde{d}([x], [y])\}$ is nonempty. So let $a' \in D_x$ and set $r := d_i(x, \phi_i(a'))$. Then we have for all $a \in D_x$

$$r + \tilde{d}([x], [y]) \geq d_i(x, \phi_i(a')) + d_i(x, \phi_i(a)) \geq d_i(\phi_i(a), \phi_i(a')) = d_A(a, a'). \quad (1.4.10)$$

In particular, setting $s := r + \tilde{d}([x], [y])$, we have $D_x \subseteq \bar{B}_s(a')$, which is a compact subset of A since A is proper. Now define $d_i^x : X_i \rightarrow \mathbb{R}$, $d_i^x(y) = d_i(x, y)$, i.e. d_i^x is the distance from the fixed point $x \in X_i$. Analogously define d_j^y . Then define

$$d_{x,y} : A \rightarrow \mathbb{R}, \quad d_{x,y}(a) := d_i^x \circ \phi_i(a) + d_j^y \circ \phi_j(a), \quad (1.4.11)$$

which is a continuous function since the sum and composition of continuous functions again results in a continuous function. But then $d_{x,y}(\bar{B}_s(a'))$ is a compact subset of $\mathbb{R}$. In particular, it attains its infimum. This means

$$\begin{aligned}\min d_{x,y}(\bar{B}_s(a')) &= \inf d_{x,y}(\bar{B}_s(a')) = \inf_{a \in \bar{B}_s(a')} \{d_i(x, \phi_i(a)) + d_j(\phi_j(a), y)\} \\ &= \inf_{a \in A} \{d_i(x, \phi_i(a)) + d_j(\phi_j(a), y)\} = \tilde{d}([x], [y]),\end{aligned}$$

where we have yet to argue for the second last equality. Since $\bar{B}_s(a') \subseteq A$ we always have $\inf d_{x,y}(\bar{B}_s(a')) \geq \inf_{a \in A} \{d_i(x, \phi_i(a)) + d_j(\phi_j(a), y)\}$. Now $\tilde{d}([x], [y]) = \inf_{a \in A} \{d_i(x, \phi_i(a)) + d_j(\phi_j(a), y)\}$ means that for all $\varepsilon > 0$ there is an $a \in A$ such that

$$\tilde{d}([x], [y]) + \varepsilon > d_i(x, \phi_i(a)) + d_j(\phi_j(a), y). \quad (1.4.12)$$

Clearly, since A is closed, we have $d(A_j, y) = \inf\{d_j(b, y) \mid b \in A_j\} = \inf\{d_j(\phi_j(a), y) \mid a \in A\} > 0$. Thus, if $a \notin \bar{B}_s(a') \supseteq D_x$ then

$$d_i(x, \phi_i(a)) + d_j(\phi_j(a), y) > \tilde{d}([x], [y]) + d(A_j, y) > \tilde{d}([x], [y]). \quad (1.4.13)$$

In other words, the a 's outside of $\bar{B}_s(a')$ are uniformly bounded away from the infimum. Setting $\varepsilon := \frac{d(A_j, y)}{2}$, the a satisfying the inequality (1.4.12) must be from $\bar{B}_s(a')$ and consequently the infima agree.

In summary, we have shown that there always exists an $a \in A$ such that $\tilde{d}([x], [y]) = d_i(x, \phi_i(a)) + d_j(\phi_j(a), y)$. But then clearly, the concatenation of the (geodesics corresponding to the) segments $[x, \phi_i(a)]$ and $[\phi_j(a), y]$ is a path of length $\tilde{d}([x], [y])$ between $[x]$ and $[y]$, making X a geodesic space.

(iv) Let $x \in X_i, y \in X_j$ and let $\varepsilon > 0$ and $a \in A$ such that $\tilde{d}([x], [y]) + \varepsilon > d_i(x, \phi_i(a)) + d_j(\phi_j(a), y)$. Since each (X_i, d_i) is a length space there is a path c_i in X_i connecting x and $\phi_i(a)$ with $L(c_i) < d_i(x, \phi_i(a)) + \frac{\varepsilon}{2}$. Analogously, there is a path c_j in X_j connecting $\phi_j(a)$ and y with $L(c_j) < d_j(\phi_j(a), y) + \frac{\varepsilon}{2}$. Then the (projection of the) concatenation $c := c_i \cdot c_j$ is a path with

$L(c) = L(c_i) + L(c_j) < d_i(x, \phi_i(a)) + d_j(\phi_j(a), y) + \varepsilon < \tilde{d}([x], [y]) + 2\varepsilon$. Thus, X is a length space.

Now suppose $\hat{d}$ is another length metric on X with $\hat{d}|_{X_i \times X_i} = d_i$ for all $i \in I$. Clearly, a path connecting $x \in X_i$ with $y \in X_j$ has to go through A . By the triangle inequality and our arguments in (i), there is exactly one point where the path will “jump” from A_i to A_j . From x to A_i and from A_j to y we already have paths of suitable length. Thus, we look for the infimum over all concatenations of such paths. But this is precisely the definition of $\tilde{d}$. $\square$

2 A first look at curvature of length spaces

After developing the basic machinery of length spaces we are ready to take a first look at curvature bounds. From now on, we assume all length spaces to be connected and geodesic. This has the following reasons: Recall that in length spaces points from two different connected components have infinite distance from each other (there is no continuous path between them). Later, we want to make statements about arbitrary triangles, i.e. triangles between any three distinct points, and for this we would like to exclude triangles with infinite sidelengths. The geodesic completeness ensures that for any triple of points the geodesic segments needed for a triangle do actually exist.

2.1 Introduction: curvature comparison in the plane

Before moving to more general curvature bounds we briefly introduce curvature comparison in the Euclidean plane. Also, it is important to emphasize that we can only talk about curvature bounds. Length spaces lack the structure that allows us to assign an explicit curvature value to them. Here, the contents are mainly taken from [BBI01, Ch. 4].

2.1.1 Definition (Non-negatively respectively non-positively curved space). *Let X be a connected geodesic space.*

- (i) *X is said to be non-positively curved if for all points in X there is a neighbourhood $U \subseteq X$ such that for all triangles $\Delta([x, y], [y, z], [z, x]) \in U$ and all $v \in [x, y]$ the following holds: Let $\bar{v} \in [\bar{x}, \bar{y}]$ be the “corresponding comparison point”, i.e. the unique point such that $d(x, v) = d(\bar{x}, \bar{v})$. Then*


$$d(v, z) \leq d(\bar{v}, \bar{z}). \quad (2.1.1)$$

- (ii) *If in the same situation as above we instead have*

$$d(v, z) \geq d(\bar{v}, \bar{z}), \quad (2.1.2)$$

for all $v \in [x, y]$ then X is said to be non-negatively curved.

Let us try to gain some intuition for this definition. In the non-positive case, any point on the side opposing a vertex has to be closer to this vertex than in the Euclidean case. This means that the side has to be in some sense “curved towards” the vertex it opposes. Thus, the condition says that triangles look thinner respectively fatter in the non-negative respectively non-positive case than the corresponding Euclidean triangles.



2.2 The model spaces

The section about the model spaces follows the contents of [BH99, Ch. 1.2] and [Hol13, Ch. 2]. We would like to not only compare length spaces to Euclidean space but also to other - already curved - spaces. In the end, we would like to say not only that a length space is at most flat respectively at least flat, but also that its curvature is bounded above respectively below by k , where k is an arbitrary real number. Formally, this comparison space - called model space - will be the unique simply connected 2-dimensional complete Riemannian manifold of constant sectional curvature k . We will however define these spaces explicitly here and rarely make use of their Riemannian structure and rather focus on their properties as metric spaces. For now we will focus on the three special cases of $k = -1, 0, 1$. We will obtain the model spaces for arbitrary k simply by rescaling. Note also that these model spaces, which we will denote by M_k^n exist in any dimension but we will most of the time use the 2-dimensional ones. Our main goal in this section will be to establish the existence of unique (up to isometry) comparison triangles in M_k^2 . To this end we will look at geodesics and other geometric objects in the model spaces such as hyperplanes explicitly.

2.2.1 Euclidean space

The simplest model space is Euclidean space. Specifically, $M_0^n := (\mathbb{R}^n, d)$, where d is the usual Euclidean metric on $\mathbb{R}^n$, i.e. $d(x, y) := \sqrt{\sum_{i=1}^n (x_i - y_i)^2}$, where $x = (x_1, \dots, x_n), y = (y_1, \dots, y_n) \in \mathbb{R}^n$. It is a well known fact that this metric is induced by the Euclidean norm which in turn is induced by the standard Euclidean scalar product. A hyperplane H in $\mathbb{R}^n$ is an affine subspace of dimension $n - 1$, i.e. any set of the form $\{v + p \mid v \in V\}$, where V is an $n - 1$ -dimensional (linear) subspace of $\mathbb{R}^n$ and $p \in \mathbb{R}^n$ an arbitrary point. In other words, a hyperplane is a linear subspace (of dimension $n - 1$) translated by a point. A vector u is orthogonal to H if it is orthogonal to the “underlying” linear subspace. Thus, it is clear that any hyperplane has (up to sign) a unique unit normal vector. If $p \in H$ and u a unit normal vector to H , then H can be written as

$$\{q \in \mathbb{R}^n \mid \langle p - q, u \rangle = 0\}. \quad (2.2.1)$$

There is another equivalent description for hyperplanes: given any two distinct points $p, q \in \mathbb{R}^n$, the set of points equidistant from p and q forms a hyperplane. It contains the midpoint of the geodesic segment $[p, q]$ and is orthogonal to $p - q$. This hyperplane is called the hyperplane bisector of p and q . Every hyperplane can be written in this form. To each hyperplane there is an associated isometry r_H called the reflection through H . r_H has the form $r_H(q) = q - 2\langle q - p, u \rangle u$, where $p \in H$ and u a unit normal vector as before. Clearly, $r_H(q) = q$ if $q \in H$. If $q \in H$, then H is the hyperplane bisector of q and $r_H(q)$ and conversely if H is the hyperplane bisector of p and q then $r_H(p) = q$.

2.2.2 The sphere

The next model space we will look at is the sphere. We set $\mathbb{S}^n := \{x \in \mathbb{R}^{n+1} \mid \langle x, x \rangle = 1\}$. On this set we define a map $d : \mathbb{S}^n \times \mathbb{S}^n \rightarrow \mathbb{R}$ by

$$d(x, y) := \arccos(\langle x, y \rangle). \quad (2.2.2)$$

Note that this is well defined since $\langle x, y \rangle \in [-1, 1]$ for all $x, y \in \mathbb{S}^n$. Clearly, d is symmetric as well as $d(x, x) = \arccos(1) = 0$ for all $x \in \mathbb{S}^n$. Now suppose $d(x, y) = 0$, i.e. $\langle x, y \rangle = 1$. But then we infer from the Cauchy-Schwarz inequality that x and y are linearly dependent, i.e. one is a multiple of the other. Since both have length 1, we obtain $x = y$. Thus, d satisfies every property of a metric except the triangle inequality, which we will show directly after a bit of preparation. Next we study some basic geometry of the sphere: by definition, a great circle on $\mathbb{S}^n$ is obtained by intersecting $\mathbb{S}^n$ with a 2-dimensional linear subspace of $\mathbb{R}^{n+1}$. Arcs of such great circles can be parametrized in the following natural way: let $x, y \in \mathbb{S}^n$ with $\langle x, y \rangle = 0$ and $a \in [0, \pi]$. Define a path $c : [0, a] \rightarrow \mathbb{S}^n$, $c(t) := \cos(t)x + \sin(t)y$. An easy computation shows that

$$\begin{aligned} d(c(t), c(s)) &= \arccos(\langle \cos(t)x + \sin(t)y, \cos(s)x + \sin(s)y \rangle) \\ &= \dots = \arccos(\cos(t)\cos(s) + \sin(t)\sin(s)) \\ &= \arccos(\cos(t-s)) = |t-s|, \end{aligned}$$

i.e. c is parametrized by arclength. Furthermore, we see that (the image of) c is contained in the great circle that is obtained by intersecting $\mathbb{S}^n$ with the subspace spanned by x and y . We call c the great arc with initial vector y from x to $c(a)$. Clearly, if $a = \pi$ then $c(\pi) = -x$ and conversely if $d(x, y) < \pi$ then there exists a unique (minimal) great arc from x to y with initial vector $y - \langle x, y \rangle x$ rescaled to length 1.

2.2.1 Definition (Spherical angle and spherical triangles). *(i) Let c, c' be two great arcs with $c(0) = c'(0) =: p \in \mathbb{S}^n$ and initial vectors v and w , respectively. Then the spherical angle between them is $\angle_p(c, c') := \arccos(\langle v, w \rangle)$.*

(ii) A spherical triangle is a collection of three distinct points $x, y, z \in \mathbb{S}^n$ and three minimal great arcs joining them pairwise and will be denoted by $\Delta(x, y, z)$.

(iii) If c and c' are the minimal great arcs joining x with y and z respectively, then $\angle_x(c, c')$ is called the vertex angle at x .

Recall the well known statement of elementary Euclidean geometry called the law of cosines: if A, B, C are the vertices of a triangle with side lengths a, b, c and angle γ at C , then $c^2 = a^2 + b^2 - 2ab\cos(\gamma)$. We aim to generalize this statement to the sphere:

2.2.2 Proposition (Spherical law of cosines). *Let $\Delta(x, y, z)$ be a spherical triangle and set $a := d(x, y), b := d(y, z), c := d(z, x)$, i.e. these are the side lengths of the three minimal great arcs forming the sides of the spherical triangle. Then for the vertex angle γ at y we have $\cos(c) = \cos(a)\cos(b) + \sin(a)\sin(b)\cos(\gamma)$.*

Proof. Let v and w be the initial vectors of the great arcs from y to x and z , respectively, then $\gamma = \arccos(\langle v, w \rangle)$. Rewriting this as $\cos(\gamma) = \langle v, w \rangle$ and rewriting $c = \arccos(\langle z, x \rangle)$ as $\cos(c) = \langle z, x \rangle$, we compute

$$\begin{aligned}
\cos(c) &= \langle z, x \rangle = \langle \cos(b)y + \sin(b)w, \cos(a)y + \sin(a)v \rangle \\
&= \dots = \cos(a)\cos(b)\langle y, y \rangle + \sin(a)\sin(b)\langle w, v \rangle \\
&= \cos(a)\cos(b) + \sin(a)\sin(b)\cos(\gamma).
\end{aligned}$$

□

Now we can prove the triangle inequality for d and consequently show some important properties of $M_0^n := (\mathbb{S}^n, d)$ as a metric space.

2.2.3 Proposition (Metric properties of $\mathbb{S}^n$). *(i) The map d defined above satisfies the triangle inequality, i.e. for all $x, y, z \in \mathbb{S}^n$ $d(x, z) \leq d(x, y) + d(y, z)$ holds with equality if and only if y lies on the minimal great arc connecting x and z .*

(ii) $(\mathbb{S}^n, d)$ is a geodesic space with geodesics being the (minimal) great arcs between points. If $d(x, y) < \pi$ then there is a unique geodesic between them.

(iii) $B_r(x)$ is convex if $r \leq \frac{\pi}{2}$ and $\bar{B}_r(x)$ is convex if $r < \frac{\pi}{2}$.

Proof. (i) Set $a := d(x, y)$, $b := d(y, z)$ and $c := d(x, z)$ and let $\Delta(x, y, z)$ be a spherical triangle and γ the vertex angle at y . Since $\cos|_{[0, \pi]} : [0, \pi] \rightarrow [-1, 1]$ is a strictly decreasing function we have that for a and b fixed and γ increasing from 0 to π that the function

$$\gamma \mapsto \cos(a)\cos(b) + \sin(a)\sin(b)\cos(\gamma) \quad (2.2.3)$$

decreases (rewriting with the addition theorems) from $\cos(b - a)$ to $\cos(b + a)$. Thus, by the spherical law of cosines and because clearly $\text{diam}(\mathbb{S}^n) = \pi$ we have $\cos(c) \geq \cos(b + a)$ and consequently $c \leq b + a$ with equality if and only if $\gamma = \pi$ and $b + a \leq \pi$. But this exactly means that y lies on the minimal great arc from x to z .

(ii) Let $d(x, y) < \pi$ and c be the minimal great arc connecting x and y . Clearly, for any total variation of c we have by (i) that $V_\sigma(c) = d(x, y)$. Thus, c is a path of (variational) length $L_d(c) = d(x, y)$, which implies that $(\mathbb{S}^n, d)$ is a geodesic space. By construction the minimal great arc is uniquely defined if $d(x, y) < \pi$ (if $d(x, y) = \pi$ then $y = -x$ and $x - \langle -x, x \rangle(-x) = 0$ does not give a vector that can be rescaled to length 1 so there is no unique great arc between them).

(iii) Recall that (if $d(x, y) < \pi$) the minimal great arc between x and y is the shorter piece of the great circle resulting from the intersection of $\mathbb{S}^n$ with the linear subspace spanned by x and y . Hence any point of this arc can be written as a linear combination $tx + sy$ of x and y . For the coefficients we must have $t + s \geq 1$ since if $t + s = 1$ this linear combination lies on the straight line segment connecting x and y and the great arc segment is “further away from the origin”. Since $d(x, y) < \pi$ there is a point $p \in \mathbb{S}^n$ such that $x, y \in \bar{B}_r(p)$ and $r < \frac{\pi}{2}$ (e.g. the midpoint of x and y). Now $z \in \bar{B}_r(p)$ if and only if $\langle p, z \rangle \geq \cos(r)$ by definition. But $t + s \geq 1$ implies that

$$\langle p, tx + sy \rangle = t\langle p, x \rangle + s\langle p, y \rangle \geq (t + s)\cos(r) \geq \cos(r), \quad (2.2.4)$$

i.e. every point in C is again $\bar{B}_r(p)$ which shows that $\bar{B}_r(p)$ is convex. The case for the open ball is analogous. □

A hyperplane H in $\mathbb{S}^n$ results from the intersection of $\mathbb{S}^n$ with an n -dimensional linear subspace V of $\mathbb{R}^{n+1}$. If we denote by R_V the usual reflection of $\mathbb{R}^{n+1}$ through V , then the reflection through H is defined as the restriction of R_V to $\mathbb{S}^n$, i.e. if $R_V : \mathbb{R}^{n+1} \rightarrow \mathbb{R}^{n+1}$, then $r_H := R_V|_{\mathbb{S}^n} : \mathbb{S}^n \rightarrow \mathbb{S}^n$ since as a restriction of an isometry r_H is again an isometry. Similar as in $\mathbb{R}^n$, the a set of points in $\mathbb{S}^n$ equidistant to $x \neq y \in \mathbb{S}^n$ is a hyperplane called the hyperplane-bisector of x and y . Alternatively this hyperplane can be described as the intersection of $\mathbb{S}^n$ with the linear subspace of $\mathbb{R}^{n+1}$ orthogonal to $x - y$. If $x \in \mathbb{S}^n$ does not lie in H then H is the hyperplane-bisector of x and $r_H(x)$ and conversely if H is the hyperplane bisector of x and y we have $r_H(x) = y$ and vice versa.

2.2.3 Hyperbolic space

Both Euclidean space and the sphere are objects where our imagination and intuition suggest their properties. This is a bit harder for hyperbolic space. Similar to the sphere we start with $\mathbb{R}^{n+1}$ but we use a different scalar product. By $\mathbb{R}_1^{n+1}$ we will denote $n+1$ -dimensional Minkowski space, i.e. $\mathbb{R}^{n+1}$ together with the scalar product

$$\langle x, y \rangle := -x_0y_0 + x_1y_1 + \dots + x_ny_n. \quad (2.2.5)$$

We then define hyperbolic n -space as the upper sheet of a paraboloid, i.e.

$$\mathbb{H}^n := \{x \in \mathbb{R}_1^{n+1} \mid \langle x, x \rangle = -1, x_0 > 0\}. \quad (2.2.6)$$

First, we need to prove the following fact:

2.2.4 Lemma (Bounding the scalar product). *For all $x, y \in \mathbb{H}^n$ we have $\langle x, y \rangle \leq -1$ with equality if and only if $x = y$.*

Proof. First note that if $x \in \mathbb{H}^n$ then clearly $x_0 \geq 1$ with $x_0 = 1$ if and only if $x = (1, 0, \dots, 0)$ (the first coordinate is the only one to pull the expression into the negative!). With this in mind, we compute:

$$\begin{aligned} \langle x, y \rangle &= -x_0y_0 + \sum_{i=1}^n x_iy_i \\ &\leq -x_0y_0 + \left(\sum_{i=1}^n x_i^2\right)^{\frac{1}{2}} \left(\sum_{i=1}^n y_i^2\right)^{\frac{1}{2}} \\ &= -x_0y_0 + (x_0^2 - 1)^{\frac{1}{2}} (y_0^2 - 1)^{\frac{1}{2}} \leq -1, \end{aligned}$$

where the last inequality holds since $x_0, y_0 \geq 1$. The second inequality in the above computation yields equality only if $x_0 = y_0$. Having this, the first inequality basically becomes the Cauchy-Schwarz inequality (viewed for the standard inner product on $\mathbb{R}^n$ by removing the x_0 and y_0 coordinate, respectively). In order for $\langle x, y \rangle$ to be -1 this also has to be an equality, i.e. and $x' := (x_1, \dots, x_n)$ is a multiple of $y' := (y_1, \dots, y_n)$ and $\sum_{i=1}^n x_i^2 = \sum_{i=1}^n y_i^2$ (since both satisfy $\langle x, x \rangle = \langle y, y \rangle = 1$ with equal first coordinate). But then we obtain $\sum_{i=1}^n x_i^2 = \sum_{i=1}^n y_i^2 = \sum_{i=1}^n \lambda x_i^2 = \lambda^2 \sum_{i=1}^n x_i^2$, which is only possible if $\lambda = 1$, i.e. if $x = y$. $\square$

Next we define a map $d : \mathbb{H}^n \times \mathbb{H}^n \rightarrow \mathbb{R}$ by

$$d(x, y) := \operatorname{arcosh}(-\langle x, y \rangle). \quad (2.2.7)$$

Since $-\langle x, y \rangle \geq 1$ this is a well defined non-negative map. Clearly, d is symmetric and $d(x, x) = 0$. By the above lemma we also have that $d(x, y) = 0$ implies $x = y$. As for the sphere we will prove the triangle inequality with the (hyperbolic) law of cosines, which we will state after developing the basic concepts of geodesics and angles. For $x, v \in \mathbb{H}^n$ with $\langle x, v \rangle = 0$ consider the path $c : \mathbb{R} \rightarrow \mathbb{H}^n$ defined by $c(t) := \cosh(t)x + \sinh(t)v$. We observe (again with the use of addition theorems) that $d(c(t), c(s)) = |t - s|$ for all $t, s \in \mathbb{R}$, i.e. c is parametrized by arclength. We define $[x, c(t)] := \{c(s) \mid s \in [0, t]\}$ as the hyperbolic segment from x to $c(t)$. If $x \neq y \in \mathbb{H}^n$ then define u to be the unit vector in direction $y + \langle x, y \rangle x$. Then u is the unique unit vector with $\langle x, u \rangle$ such that $y = \cosh(r)x + \sinh(r)u$, where $r := d(x, y)$. Similar to the sphere we will call u the initial vector of the hyperbolic segment $[x, y]$.

2.2.5 Definition (Hyperbolic angle and triangles). *(i) Let $[x, c(t)]$ and $[x, c'(s)]$ be two hyperbolic segments both emanating from $x \in \mathbb{H}^n$ with initial vectors u and v respectively. Then the hyperbolic angle between c and c' is defined as the unique $\alpha \in [0, \pi]$ such that $\cos(\alpha) = \langle u, v \rangle$. Note that here we cannot take $\arccos$ since $\langle u, v \rangle$ may be smaller than 1.*

(ii) A hyperbolic triangle consists of three (distinct) points $x, y, z \in \mathbb{H}^n$ and the (unique!) hyperbolic segments between them and will be denoted by $\Delta(x, y, z)$.

(iii) The vertex angle at y is defined as the hyperbolic angle between $[y, x]$ and $[y, z]$.

2.2.6 Proposition (Hyperbolic law of cosines). *Let $\Delta(x, y, z)$ be a hyperbolic triangle, $a := d(x, y)$, $b := d(y, z)$, $c := d(x, z)$ and γ the vertex angle at y . Then:*

$$\cosh(c) = \cosh(a) \cosh(b) - \sinh(a) \sinh(b) \cos(\gamma). \quad (2.2.8)$$

Proof. Let u and v be the initial vectors of the hyperbolic segments $[y, x]$ and $[y, z]$ respectively. Then $\gamma = \langle u, v \rangle$. Using $\langle x, u \rangle = 0$, $\langle y, v \rangle = 0$ and $\langle y, y \rangle = -1$, we compute

$$\begin{aligned} \cosh(c) &= -\langle x, z \rangle \\ &= -\langle \cosh(a)y + \sinh(a)u, \cosh(b)z + \sinh(b)v \rangle \\ &= \dots = -\cosh(a) \cosh(b) \langle y, y \rangle - \sinh(a) \sinh(b) \langle u, v \rangle \\ &= \cosh(a) \cosh(b) - \sinh(a) \sinh(b) \cos(\gamma). \end{aligned}$$

□

Having the above at hand, we are ready to prove the triangle inequality of d and some resulting properties.

2.2.7 Proposition (Metric properties of $\mathbb{H}^n$). *(i) The map $d : \mathbb{H}^n \times \mathbb{H}^n \rightarrow \mathbb{R}$ satisfies the triangle inequality, i.e. for all $x, y \in \mathbb{H}^n$ we have*

$$d(x, z) \leq d(x, y) + d(y, z) \quad (2.2.9)$$

with equality if and only if y lies on the hyperbolic segment $[x, z]$.

(ii) $(\mathbb{H}^n, d)$ is a uniquely geodesic space with the unique geodesic joining x and y being the hyperbolic segment $[x, y]$.

(iii) The geodesic segments in $\mathbb{H}^n$ exactly correspond to the intersections of $\mathbb{H}^n$ with a 2-dimensional subspace of $\mathbb{R}_1^{n+1}$.

(iv) For all $x \in \mathbb{H}^n$ and all $r > 0$ $B_r(x)$ and $\bar{B}_r(x)$ are convex.

Proof. (i) Set $a := d(x, y)$, $b := d(y, z)$, $c := d(x, z)$ and let u, v be the initial vectors of the hyperbolic segments joining y and x, z respectively and consider the hyperbolic triangle $\Delta(x, y, z)$ with vertex angle γ at y . We argue similarly to the case of the sphere: since $\cos$ increases on the interval $[0, \pi]$ we have for fixed a, b and γ increasing from 0 to π that the function

$$\gamma \mapsto \cosh(a) \cosh(b) - \sinh(a) \sinh(b) \cos(\gamma) \quad (2.2.10)$$

increases (rewritten with the addition theorems) from $\cosh(b - a)$ to $\cosh(b + a)$. Thus, the hyperbolic law of cosines gives $\cosh(c) \geq \cosh(b + a)$ and consequently $c \geq a + b$ with equality if and only if $\gamma = \pi$. And $\gamma = \pi$ if and only if $y \in [x, z]$.

(ii) The only thing left to note is that because of $d(x, z) = d(x, y) + d(y, z)$ if and only if $y \in [x, z]$ we have for any total variation of the (restriction of the) path corresponding to the geodesic segment $[x, z]$ that $L_d(c) = d(x, z)$, implying that $(\mathbb{H}^n, d)$ is a geodesic space.

(iii) Clearly, any hyperbolic segment with starting point x and initial vector u lies in the linear subspace spanned by those two. Conversely any non-empty intersection of $\mathbb{H}^n$ with a 2-dimensional subspace gives a geodesic segment.

(iv) This works analogously to the sphere: we have that any point of $[x, y]$ can be written as $tx + sy$ with $t + s \geq 1$ and $t, s \geq 0$. Also, $x \in \bar{B}_r(z)$ if and only if $-\langle x, z \rangle \leq \cosh(r)$. So let $x, y \in \bar{B}_r(z)$ and compute

$$\begin{aligned} -\langle z, tx + sy \rangle &= -t\langle z, x \rangle - s\langle z, y \rangle \\ &\geq t \cosh(r) + s \cosh(r) \\ &\geq (t + s) \cosh(r) \geq \cosh(r). \end{aligned}$$

Thus, any point of the segment $[x, y]$ lies again in $\bar{B}_r(z)$ and so $\bar{B}_r(z)$ is convex. Again, the case for the open ball is similar. $\square$

Finally, we take a brief look at hyperplanes in $\mathbb{H}^n$: similarly to the sphere, a hyperplane in $\mathbb{H}^n$ results from an intersection of $\mathbb{H}^n$ with an n -dimensional subspace of $\mathbb{R}_1^n$ and any such hyperplane is again isometric to $\mathbb{H}^{n-1}$. For any two distinct $x, y \in \mathbb{H}^n$ the set of points equidistant from x and y form a hyperplane which is called the hyperplane bisector of x and y . The hyperplane bisector of x and y is obtained by intersecting $\mathbb{H}^n$ with the orthogonal complement of $x - y$, which is an n -dimensional linear subspace of $\mathbb{H}^n$. Again, for every hyperplane H there is an associated isometry called the reflection through H . This is given by $r_H : \mathbb{H}^n \rightarrow \mathbb{H}^n$, $x \mapsto x - 2\langle x, u \rangle u$, where u is the (up to sign unique) unit vector orthogonal to the n -dimensional subspace that gives H when intersected with $\mathbb{H}^n$. At last, we remark as in the case of Euclidean space and the sphere that the set of fixed points of r_H is H . If $x \notin H$ then H is the hyperplane bisector of x and $r_H(x)$ and conversely, if H is the hyperplane bisector of x and y then $r_H(x) = y$ and vice versa.

2.2.4 The general model space M_k^n

In this last subsection we briefly discuss how to obtain the general model space M_k^n from rescaling the above mentioned constructions. Let us first look at the sphere, i.e. $k > 0$: since in no way did the choice of radius of our sphere matter, we can take as an underlying set a bigger

respectively smaller sphere and leave the distance function unchanged. More precisely, we define $S_k^n := \{x \in \mathbb{R}^{n+1} \mid \langle x, x \rangle = \frac{1}{k}\}$ and set $M_k^n := (S_k^n, d)$, where d is metric we explicitly defined on the sphere. Alternatively, instead of changing the set one can also rescale the distance function. By setting $d_k(x, y) := \frac{1}{\sqrt{k}} \arccos(\langle x, y \rangle)$, we define $M_k^n := (\mathbb{S}^n, d_k)$. Clearly, these two definitions are equivalent.

The same can be done for hyperbolic space. Either we rescale the equation determining our underlying set or we rescale the distance function. Of course, the condition $x_0 > 0$ still applies if we change the set.

2.2.8 Definition (Diameter of M_k^2). *We set $D_k := \text{diam}(M_k^2)$. Thus, $D_k = \infty$ if $k \leq 0$ and $D_k = \frac{\pi}{\sqrt{k}}$ if $k > 0$.*

Since in particular the law of cosines is not influenced by the number in the defining equations of $\mathbb{S}^n$ respectively $\mathbb{H}^n$, we obtain the following corollary.

2.2.9 Corollary (Generalized law of cosines). *Let $\Delta(x, y, z)$ be a (geodesic) triangle in M_k^n , $a := d(x, y)$, $b := d(y, z)$, $c := d(z, x)$ and γ the vertex angle at y .*

$$\begin{aligned} \text{If } k = 0 : & \quad c^2 = a^2 + b^2 - 2ab \cos(\gamma), \\ \text{if } k < 0 : & \quad \cosh(\sqrt{-k}c) = \cosh(\sqrt{-k}a) \cosh(\sqrt{-k}b) - \sinh(\sqrt{-k}a) \sinh(\sqrt{-k}b) \cos(\gamma), \\ \text{and if } k > 0 : & \quad \cos(\sqrt{k}c) = \cos(\sqrt{k}a) \cos(\sqrt{k}b) - \sin(\sqrt{k}a) \sin(\sqrt{k}b) \cos(\gamma). \end{aligned}$$

□

In addition, the following metric properties are preserved. The argument is the same as for the law of cosines: the numbers in the defining equations are irrelevant.

2.2.10 Corollary (Metric properties of M_k^2). *(i) M_k^2 is a geodesic space. Furthermore, if $d(x, y) < D_k$ then the connecting geodesic is unique.*

(ii) $B_r(x)$ is convex if $r \leq \frac{D_k}{2}$ and $\bar{B}_r(x)$ is convex if $r < \frac{D_k}{2}$.

□

2.3 Uniqueness of comparison triangles and Alexandrov's lemma

We are now in a position to show that comparison triangles are unique in any 2-dimensional model space. For this section we again will mainly follow [BH99, Ch. I.2] and [Hol13, Ch. 2].

First we would like to rule out a potential problem concerning angles: on the one hand, we defined the spherical respectively hyperbolic angle for our model spaces explicitly. On the other hand, one can also use the Alexandrov angle since the model spaces are length spaces. Fortunately, these notions coincide. Since the proof is rather long and technical, we will just give a reference, see e.g. [BH99, Ch. I, Prop. 2.9].

2.3.1 Proposition (Equivalence of angle definitions). *Let $[x, y]$ and $[x, z]$ be two geodesic segments in M_k^n emanating from the same point $x \in M_k^n$. Then the spherical respectively hyperbolic angle (depending on k) is equal to the Alexandrov angle between them.* □

Thus, we have a well defined notion of angles in any model space. For better readability, we will (analogously to our first definition of a comparison angle in the Euclidean plane) denote the angle between two geodesic segments $[x, y], [x, z]$ in M_k^n as $\angle_x(y, z)$, i.e. omitting the paths to which these segments correspond. This is well defined since geodesics are unique in any model space except when $d(x, y) = D_k$, in which case we must specify which segment we consider! If the context requires us to specify in which model space an angle is taken then we will denote the angle by $\angle_x^{(k)}(y, z)$.

For the uniqueness of comparison triangles we prove the following fact about the isometry group of M_k^n .

2.3.2 Proposition (Transitivity of $\text{Isom}(M_k^n)$). *Let $x_1, \dots, x_n$ and $y_1, \dots, y_n$ be pairwise equidistant points in M_k^n , i.e. $d(x_i, x_j) = d(y_i, y_j)$ for all $i, j = 1, \dots, n$. Then there exists an isometry $\phi \in \text{Isom}(M_k^n)$ such that $\phi(x_i) = y_i$ for all $i = 1, \dots, n$. Furthermore, this isometry can be obtained by composing at most n reflections through hyperplanes.*

Proof. We prove the statement by induction. Let $n = 1$, then by construction we have that the points equidistant from x_1 and y_1 form a hyperplane H , namely the hyperplane bisector of x_1 and y_1 . And by construction we have for the hyperplane bisector of x_1 and y_1 that $r_H(x_1) = y_1$. Now let $x_1, \dots, x_n$ and $y_1, \dots, y_n$ as in the assumption and suppose that the statement holds for $n - 1$ points. By induction hypothesis there is an isometry $\phi \in \text{Isom}(M_k^n)$ such that $\phi(x_i) = y_i$ for $i = 1, \dots, n - 1$ and ϕ is a composition of at most $n - 1$ reflections. If $\phi(x_n) \neq y_n$ (otherwise we are done) let H be the hyperplane bisector of $\phi(x_n)$ and y_n and denote by r_H the reflection through H . Observe that $y_i \in H$ for $i = 1, \dots, n - 1$ since $d(y_i, \phi(x_n)) = d(\phi(x_i), \phi(x_n)) = d(x_i, x_n) = d(y_i, y_n)$. Thus, $r_H(y_i) = y_i$ for $i = 1, \dots, n - 1$ and so $r_H \circ \phi$ is an isometry with the desired property. $\square$

From now on, if we work with triangles in general length spaces and triangles in the model spaces simultaneously, we will call any triangle in a model space a comparison triangle (for a suitable triangle from a general length space) and denote it by $\bar{\Delta}(\bar{x}, \bar{y}, \bar{z})$. If we deal with triangles exclusively in model spaces we may omit the bar and write Δ instead of $\bar{\Delta}$.


2.3.3 Theorem (Existence of comparison triangles in M_k^2). *Let X be a length space, $k \in \mathbb{R}$ and $x, y, z \in X$ such that $d(x, y) + d(y, z) + d(z, x) < 2D_k$. Then there exist $\bar{x}, \bar{y}, \bar{z} \in M_k^2$ such that $d(x, y) = d(\bar{x}, \bar{y})$, $d(y, z) = d(\bar{y}, \bar{z})$ and $d(z, x) = d(\bar{z}, \bar{x})$. In particular, for a triangle $\Delta([x, y], [y, z], [z, x])$ in X there is a unique (up to isometry) comparison triangle $\bar{\Delta}(\bar{x}, \bar{y}, \bar{z})$ in M_k^2 .*

Proof. Set $a := d(x, y)$, $b := d(y, z)$ and $c := d(z, x)$. Without loss of generality we assume $a \leq b \leq c$ (otherwise rename the points). By the triangle inequality we obtain $c \leq a + b$. Thus, if $k > 0$ we have $c < \frac{\pi}{\sqrt{k}}$. This means that we can use a, b and c in the appropriate law of cosines in 2.2.9 to obtain a unique $\gamma \in [0, \pi]$ solving the equation. Now fix any $\bar{y} \in M_k^2$ and construct two geodesic segments $[\bar{y}, \bar{x}]$ and $[\bar{y}, \bar{z}]$ of length a and b , respectively, such that the angle between them is γ . Then the law of cosines gives $d(\bar{x}, \bar{z}) = c$.

The uniqueness of the comparison triangle is a special case of 2.3.2: two comparison triangles for the same triangle in X are given respectively by a collection of three points of pairwise equal distance. Thus, we find an isometry mapping the vertices onto each other. $\square$

Next we will prove Alexandrov's lemma, which is an elementary yet very useful statement. While the assumptions are quite technical, there is a simple and almost beautiful intuition behind it.

The first assertion seems almost trivial (at least in the Euclidean plane): it states that in a concave quadrangle the sum of the lengths of the two sides adjacent to the reflex angle is always shorter than the sum of the lengths of the other two sides. For the second assertion, consider two triangles in the Euclidean plane sharing an edge. Suppose the sum of the interior angles at the top end of this edge is at least π . Denote the vertex at this angle by C and the other endpoint of this edge by A . Now remove the common edge and imagine the remaining edges as iron bars which are joined by hinges at the vertices (so that the length of the sides remains unchanged). Alexandrov's lemma basically states that when one decreases the interior angle at C to π , i.e. "straightens" the top bar then all the other interior angles increase. In particular, this causes A to move further away from C . Since the proof does not use more than the law of cosines, the result holds in any model space.



We will make use of the following notation: let $x, y \in M_k^2$ be two distinct points with $d(x, y) < D_k$. Then there is a unique "geodesic line" $\bar{c} : \mathbb{R} \rightarrow M_k^2$ extending the geodesic $c : [0, l] \rightarrow M_k^2$ corresponding to the geodesic segment $[x, y]$. $\bar{c}$ splits M_k^2 into two connected components. We say two points lie on opposite sides of this line if they lie in different connected components of $M_k^2 \setminus \{\bar{c}(t) \mid t \in \mathbb{R}\}$.

2.3.4 Lemma (Alexandrov's lemma). *Let A, B, C and D be four distinct points in M_k^2 such that the circumference of the quadrangle they are forming is less than $2D_k$, i.e. $d(A, B) + d(B, C) + d(C, D) + d(D, A) < 2D_k$ and that B and D lie on different sides of the geodesic line through A and C . Consider the triangles $\Delta(A, B, C)$ and $\Delta(A, C, D)$. Let α_1, β and γ_1 be the interior angles in $\Delta(A, B, C)$ at the vertices A, B and C , respectively. Let α_2, γ_2 and δ be the interior angles in $\Delta(A, C, D)$ at the vertices A, C and D , respectively. Suppose $\gamma_1 + \gamma_2 \geq \pi$. Then*

(i) $d(B, C) + d(C, D) \leq d(A, B) + d(A, D)$.

(ii) *Suppose additionally that $\Delta(A', B', D')$ is a triangle in M_k^2 such that $d(A', B') = d(A, B)$, $d(A', D') = d(A, D)$ and $d(B', D') = d(B, C) + d(C, D)$. Let C' be the unique point on $[B', D']$ such that $d(B', C') = d(B, C)$ (and consequently $d(C', D') = d(C, D)$). Denote by α', β' and δ' the interior angles of $\Delta(A', B', D')$ at A', B' and D' , respectively. Then $\alpha' \geq \alpha_1 + \alpha_2, \beta' \geq \beta, \delta' \geq \delta$ and $d(A', C') \geq d(A, C)$. Any of these is an equality if and only if all are. This happens if and only if $\gamma_1 + \gamma_2 = \pi$.*

Proof. (i) Let $\tilde{B}$ be the unique point in M_k^2 with $d(\tilde{B}, C) = d(D, C)$ such that C lies on the geodesic segment $[\tilde{B}, B]$. Equivalently, one can describe this point as follows: The intersection

of the geodesic line through B and C with the circle of radius $d(D, C)$ with midpoint C yields two points. $\tilde{B}$ is the one such that C lies between B and $\tilde{B}$. Since $\gamma_1 + \gamma_2 \geq \pi$, the angle at C between $[C, A]$ and $[C, \tilde{B}]$ is not greater than the angle at C between $[C, A]$ and $[C, D]$. While this seems trivial when considering the plane, in general it can be argued as follows: $\gamma_1 + \gamma_2 \geq \pi$ implies that the points A and D lie on opposite sides of the line through B and $\tilde{B}$ and so the segment $[C, \tilde{B}]$ lies “between” the segments $[C, A]$ and $[C, D]$. More precisely, let $c_A, c_{\tilde{B}}$ and c_D be the geodesics corresponding to the segments $[C, A]$, $[C, \tilde{B}]$ and $[C, D]$, respectively. Then for small enough parameters $t > 0$ we have $d(c_A(t), c_{\tilde{B}}(t)) \leq d(c_A(t), c_D(t))$. Now the explicit description of the angle (1.2.7) between them implies that $\bar{\angle}_C(A, \tilde{B}) \leq \bar{\angle}_C(A, D)$ (since arccos is decreasing). Recall the consequence of the law of cosines stating that when keeping two side lengths fixed, the third side increases as the vertex angle it opposes increases. This in particular implies $d(A, \tilde{B}) \leq d(A, D)$, as the two corresponding triangles $\Delta(A, C, D)$ and $\Delta(A, C, \tilde{B})$ have two sides of (pairwise) equal length since by construction $d(C, D) = d(C, \tilde{B})$. But then

$$\begin{aligned} d(A, B) + d(A, D) &\geq d(A, B) + d(A, \tilde{B}) \\ &\geq d(B, \tilde{B}) = d(B, C) + d(C, \tilde{B}), \end{aligned}$$

where the last equality is because C lies on the geodesic segment $[B, \tilde{B}]$.

(ii) Since $d(A, D) = d(A', D')$, by (i) we have $d(A', D') \geq d(A, \tilde{B})$. By the triangle inequality, we have $d(B', D') = d(B, C) + d(C, D) \geq d(B, D)$. In both cases we have equality if and only if $C \in [B, D]$, i.e. if $\gamma_1 + \gamma_2 = \pi$. The next argument is again similar to (i): Since by construction $d(A, B) = d(A', B')$ and $d(A, D) = d(A', D')$, the triangles $\Delta(A, B, D)$ and $\Delta(A', B', D')$ have two sides of equal length. Since $d(B, D) \leq d(B', D')$, the third side in $\Delta(A, B, D)$ is shorter than in $\Delta(A', B', D')$ and hence the corresponding angle is smaller, i.e. $\alpha_1 + \alpha_2 \leq \alpha'$. Similarly, by construction we have $d(A, B) = d(A', B')$ and $d(B, \tilde{B}) = d(B', D')$ and so the triangles $\Delta(A, B, \tilde{B})$ and $\Delta(A', B', D')$ have two sides of equal length. Since $d(A, \tilde{B}) \leq d(A', D')$, the third side in $\Delta(A, B, \tilde{B})$ is shorter than in $\Delta(A', B', D')$ and so $\beta \leq \beta'$. Replacing B by D the exact same argument yields $\delta \leq \delta'$. Now consider the triangles $\Delta(A, B, C)$ and $\Delta(A', B', C')$ which again have two sides of equal length. Since the angle opposing the third side is greater in $\Delta(A', B', C')$, its third side must be longer as well, i.e. $d(A', C') \geq d(A, C)$. Finally, we note that all the above inequalities become equalities if and only if $\gamma_1 + \gamma_2 = \pi$, in which case $\Delta(A', B', D')$ is isometric to the union of $\Delta(A, B, C)$ and $\Delta(A, C, D)$ with the common edge $[A, C]$ deleted, meaning that equality holds everywhere. $\square$

2.3.5 Remark (Other direction of Alexandrov’s lemma). Alexandrov’s lemma also holds when the assumptions are reversed: loosely speaking, if the angle at C is smaller than π then this straightening process causes A to move closer to C and so on. Since we are mainly interested in curvature bounds from above, the case we treated in 2.3.4 suffices.

3 CAT(k) spaces

Now we have all tools necessary at hand to discuss geodesic spaces with curvature bounded above. The idea will be the same as introduced in 2.1 and intuitively amounts to saying that triangles in the space look thinner than the triangles in some (already curved) model space. The quantity k indicates with which M_k^2 we compare and tells in some sense how “strong” the curvature bound is. Thus, every CAT(k) space is also a CAT(l) space for all $l > k$ as we will prove later.

Before we continue, we would like to modify the statement we made at the beginning of chapter 2. We announced that all length spaces are required to be geodesic and connected. This assumption is too strong in some cases: if $k > 0$, we actually do not care how points in X further apart than D_k behave, since this cannot be captured by comparison with a (bounded) sphere. Additionally, the $\text{CAT}(k)$ inequality is in essence a local property. Therefore, we require a $\text{CAT}(k)$ space only to be D_k -geodesic. Clearly, if $k \leq 0$ then being D_k -geodesic is equivalent with being geodesic. Also note that a $\text{CAT}(k)$ space may be disconnected, although local path connectedness is ensured in any length space, cf. [BBI01, Exercise 2.1.9].

3.1 Basic properties

We now put these thoughts into a concrete definition. We will for the most part follow [BH99, Ch. II.1], [Hol13, Ch. 3], [Jos97, Ch. 2], [Shi93, Ch. 3] and [AKP19a, Ch. 8].

3.1.1 Definition (The $\text{CAT}(k)$ inequality). *Let X be a length space, $k \in \mathbb{R}$ and $\Delta(x, y, z)$ a triangle in X with $d(x, y) + d(y, z) + d(z, x) < 2D_k$. Let $\bar{\Delta}(\bar{x}, \bar{y}, \bar{z})$ be a comparison triangle for $\Delta(x, y, z)$ in M_k^2 . Then $\Delta(x, y, z)$ is said to satisfy the $\text{CAT}(k)$ inequality if the following holds for all points in $\Delta(x, y, z)$: Let $p, q \in \Delta(x, y, z)$ and let $\bar{p}, \bar{q} \in \bar{\Delta}(\bar{x}, \bar{y}, \bar{z})$ be comparison points for p and q , respectively. Recall that this means if, say $p \in [x, y]$, then $\bar{p} \in [\bar{x}, \bar{y}]$ and $d(x, p) = d(\bar{x}, \bar{p})$. Then*

$$d(p, q) \leq d(\bar{p}, \bar{q}). \quad (3.1.1)$$

While the formulation is slightly different than the one mentioned in the non-positively curved case in 2.1.1, it basically says the same: instead of saying that all points on a side of a triangle in X have to be closer to the opposing vertex (compared to a Euclidean comparison triangle), we now require that all pairs of points are closer to each other than in a comparison triangle in the appropriate model space. Clearly, both conditions make the triangles in X look thinner than the ones we compare them to. Note that the $\text{CAT}(k)$ inequality is trivially satisfied if p and q belong to the same side. Say $p, q \in [x, y]$ and assume without loss of generality that p is closer to x than q . Since $[x, y]$ is isometric to $[\bar{x}, \bar{y}]$ and both are geodesic segments we consequently have $d(p, q) = d(x, y) - d(x, p) - d(q, y) = d(\bar{x}, \bar{y}) - d(\bar{x}, \bar{p}) - d(\bar{q}, \bar{y}) = d(\bar{p}, \bar{q})$.

3.1.2 Definition (Circumference). *Let $\Delta([x, y], [y, z], [z, x])$ be a triangle in a length space X , $k \in \mathbb{R}$. The circumference of the triangle $\Delta([x, y], [y, z], [z, x])$ is defined as the sum $d(x, y) + d(y, z) + d(z, x)$.*

3.1.3 Definition ($\text{CAT}(k)$ and local $\text{CAT}(k)$ space). *Let $k \in \mathbb{R}$ and let X be a length space.*

- (i) *X is called a $\text{CAT}(k)$ space if it is D_k -geodesic and all triangles in X with circumference $< 2D_k$ satisfy the $\text{CAT}(k)$ inequality.*
- (ii) *X is called a local $\text{CAT}(k)$ space if for all $x \in X$ there exists $r > 0$ such that $B_r(x)$ with the induced length metric (cf. 1.3.2) is a $\text{CAT}(k)$ space.*

From the definition it is clear that the bound on the circumference is only relevant when $k > 0$. On any sphere, triangles that are too large (with respect to the circumference) cannot have geodesics as sidelengths. Note that any space that is non-positively curved in the sense of 2.1.1 is locally a $\text{CAT}(0)$ space.

Next, we will show some elementary properties of $\text{CAT}(k)$ spaces. The first claim in the proposition below in part deals with the continuous variation of geodesic segments. What we mean by this precisely is the following: let x and y be points in a $\text{CAT}(k)$ space X and let $x_n \rightarrow x$

and $y_n \rightarrow y$ be converging sequences. Let γ be the curve corresponding to the geodesic segment $[x, y]$ and let γ_n be the curve corresponding to the geodesic segment $[x_n, y_n]$ (note that since we will also show uniqueness of geodesics, it is okay to talk about *the* geodesic segments between points). We assume these curves to be linearly parameterized, i.e. having $[0, 1]$ as domain. Technically, these curves cease to be geodesics with respect to our definition. However, the image clearly remains the same and in this case we can compare the curves easily. Namely, we require that the curves γ_n converge uniformly to γ . Before continuing with the proposition we prove a useful lemma which bounds the spread of geodesics with respect to the distance of its endpoints in a model space. Unfortunately, we will for the first time make use of the strong differential machinery available in the model spaces.

3.1.4 Lemma (Bounding the spread of geodesics). *Let $k > 0$. Then there exists a function $C : [0, D_k) \rightarrow \mathbb{R}$ such that $\lim_{r \rightarrow 0} C(r) = 1$ and for all $z \in M_k^2$ and all $x, y \in B_r(z)$ we have*

$$d(x_s, y_s) \leq sC(r)d(x, y), \quad (3.1.2)$$

where $x_s \in [z, x]$ is the unique point such that $d(z, x_s) = sd(z, x)$, $s \in (0, 1)$.

Proof. We view $B_r(z) \subseteq M_k^2$ as a ball with radius r in $(\mathbb{R}^2, g_k)$, where the Riemannian metric g_k is expressed by

$$ds^2 = d\rho^2 + \frac{1}{k} \sin(\sqrt{k}\rho)^2 d\theta^2 \quad (3.1.3)$$

in polar coordinates (ρ, θ) . We will denote the associated (distance) metric on $B_r(z)$ by d_k . Next, we observe that x_s retains its defining properties also from the Euclidean point of view, i.e. it still is a point on the Euclidean geodesic from z to x with distance $sd(z, x)$ from z . The same clearly holds for y_s so in particular we have $d_{\mathbb{R}^2}(x_s, y_s) = sd_{\mathbb{R}^2}(x, y)$, where $d_{\mathbb{R}^2}$ denotes the Euclidean metric (think of the intercept theorem). Since $B_r(z)$ is compact (and equals the closure of $B_r(z)$), we infer the existence of a constant $C(r) \searrow 1$ as $r \searrow 0$ that relates the Euclidean norm of a tangent vector v to the norm induced by g_k via

$$\frac{1}{\sqrt{C(r)}} \|v\|_k \leq \|v\|_{\mathbb{R}^2} \leq \sqrt{C(r)} \|v\|_k. \quad (3.1.4)$$

More generally, such an estimate holds for any Riemannian metric g on $\mathbb{R}^n$. Choose Riemannian normal coordinates around $p \in \mathbb{R}^n$. The continuity of g implies $g_q(v, v) \rightarrow 1$ as $r \rightarrow 0$ for all $q \in B_r(p)$, $v \in \mathbb{S}^{n-1}$, where $\mathbb{S}^{n-1}$ is the Euclidean unit sphere. To see this, first note that clearly $g_p(v, v) = \langle v, v \rangle = 1$, where $\langle \cdot, \cdot \rangle$ denotes the Euclidean scalar product. Now writing $g_q(v, v)$ as $v^\top G(q)v$, we compute

$$\begin{aligned} |v^\top G(q)v - v^\top G(p)v| &= |v^\top (G(q) - \mathbb{I}_n)v| \\ &\leq \|v^\top\| \|G(q) - \mathbb{I}_n\| \|v\| = \|G(q) - \mathbb{I}_n\|, \end{aligned}$$

which goes to 0 as $r \rightarrow 0$ since then necessarily also $q \rightarrow p$ and so $G(q) \rightarrow G(p) = \mathbb{I}_n$. Setting

$$C_1(r) := \sup_{\substack{q \in B_r(p) \\ v \in \mathbb{S}^{n-1}}} g_q(v, v), \quad \frac{1}{C_2(r)} := \inf_{\substack{q \in B_r(p) \\ v \in \mathbb{S}^{n-1}}} g_q(v, v) \text{ and } C(r) := \max\{C_1(r), C_2(r)\}, \quad (3.1.5)$$

the claim follows for all (Euclidean) unit vectors. Finally, we observe that this holds for arbitrary v as well by considering $\frac{1}{\|v\|}v \in \mathbb{S}^{n-1}$ and cancelling the scalars in the estimate appropriately.

This equivalence of norms in particular holds for the tangent vectors of geodesics, and so integrating leads to a corresponding estimate for the length of curves. Since the distance between two points is given by the length of a shortest curve, we arrive at

$$\frac{1}{\sqrt{C(r)}}d_k(x, y) \leq d_{\mathbb{R}^2}(x, y) \leq \sqrt{C(r)}d_k(x, y) \quad (3.1.6)$$

for all $x, y \in B_r(z)$. Thus, we obtain $d_k(x_s, y_s) \leq \sqrt{C(r)}d_{\mathbb{R}^2}(x_s, y_s) = \sqrt{C(r)}sd_{\mathbb{R}^2}(x, y) \leq sC(r)d_k(x, y)$. $\square$

Note that because of the above generalization, we obtain a similar constant also in hyperbolic space.


3.1.5 Proposition (Basic properties of $\text{CAT}(k)$ spaces). *Let X be a $\text{CAT}(k)$ space.*

- (i) *For all $x, y \in X$ with $d(x, y) < D_k$ there exists a unique geodesic joining them. Moreover, this geodesic varies continuously with its endpoints.*
- (ii) *If $\gamma : [0, l] \rightarrow X$ is a local geodesic with $L(\gamma) < D_k$, then it is already a geodesic.*
- (iii) *$B_r(x)$ and $\bar{B}_r(x)$ are convex for all $x \in X$ and all $r \leq \frac{D_k}{2}$ respectively $r < \frac{D_k}{2}$.*

Proof. (i) Let $d(x, z) < D_k$ and suppose $[x, z]$ and $[x, z]'$ are two geodesic segments joining x and z . Let $y \in [x, z]$ and $y' \in [x, z]'$ be such that they have equal distance from x , i.e. $d(x, y) = d(x, y')$. Consider the division of $[x, z]$ into the two segments $[x, y]$ and $[y, z]$. Clearly, any comparison triangle in M_k^2 for the triangle $\Delta([x, y], [y, z], [x, z]')$ is degenerate. Moreover, we have that the comparison points for y and y' are the same, i.e. $\bar{y} = \bar{y}'$. Thus, the $\text{CAT}(k)$ inequality implies $0 = d(\bar{y}, \bar{y}') \geq d(y, y') \geq 0$, and consequently $d(y, y') = 0$ and so $y = y'$. Since these points were chosen arbitrarily, the segments $[x, z]$ and $[x, z]'$ have to be equal.

Concerning the continuous variation, first note that 3.1.4 guaranties the existence of a constant C (depending on l and k) such that two linearly parameterized geodesics c and c' in M_k^2 with common starting point of length not greater than $l < D_k$ satisfy $d(c(t), c'(t)) \leq Cd(c(1), c'(1))$ for all $t \in [0, 1]$. Now given $l < D_k$, let $x_n \rightarrow x$ and $y_n \rightarrow y$ be two sequences in X with $d(x_n, y_n) < l$ and $d(x, y_n) < l$ for all n . Let γ, γ_n and γ'_n be linear parameterizations of the geodesics corresponding to the segments $[x, y], [x_n, y_n]$ and $[x, y_n]$, respectively. More precisely, let $\gamma(0) = x, \gamma(1) = y, \gamma_n(0) = x_n, \gamma_n(1) = y_n, \gamma'_n(0) = x$ and $\gamma'_n(1) = y_n$. Consider the triangles $T_1 := \Delta([x, y], [y, y_n], [y_n, x])$ and $T_2 := \Delta([x, y_n], [y_n, x_n], [x_n, x])$, having the edge $[x, y_n]$ in common. In particular, $\gamma(t) \in [x, y], \gamma_n(t) \in [x_n, y_n]$ and $\gamma'_n(t) \in [x, y_n]$ for all $t \in [0, 1]$. Let $\bar{\Delta}(\bar{x}, \bar{y}, \bar{y}_n)$ $\bar{\Delta}(\bar{x}, \bar{y}_n, \bar{x}_n)$ be comparison triangles for the above defined triangles, respectively, also sharing the edge $[\bar{x}, \bar{y}_n]$. Consider “comparison geodesics” c, c_n, c'_n corresponding to the segments $[\bar{x}, \bar{y}], [\bar{x}_n, \bar{y}_n]$ and $[\bar{x}, \bar{y}_n]$, respectively, also linearly parameterized. On the one hand, we infer by the $\text{CAT}(k)$ inequality applied to T_1 that $d(\gamma(t), \gamma'_n(t)) \leq d(c(t), c'_n(t))$. By 3.1.4 and the above considerations we have $d(c(t), c'_n(t)) \leq Cd(c(1), c'_n(1)) = Cd(y, y_n)$. On the other hand, we also obtain by the $\text{CAT}(k)$ inequality in T_2 that $d(\gamma_n(t), \gamma'_n(t)) \leq d(c_n(t), c'_n(t))$. To be reasonably rigorous, we have to mention a small technicality for applying the lemma in this case: we want to bound the distance with respect to the endpoints, but the geodesics how we defined them here “run in the wrong direction”. To correct this, we simply define $\tilde{c}_n(t) := c_n(1 - t)$ and $\tilde{c}'_n(t) := c'_n(1 - t)$. Then $\tilde{c}_n$ and $\tilde{c}'_n$ have common starting point y_n and we have

$$\begin{aligned}
d(c_n(t), c'_n(t)) &= d(\tilde{c}_n(1-t), \tilde{c}'_n(1-t)) \\
&\leq Cd(\tilde{c}_n(1), \tilde{c}'_n(1)) = Cd(c_n(0), c'_n(0)) \\
&= Cd(x_n, x).
\end{aligned}$$



By adding these two cases together and applying the triangle inequality, we finally arrive at

$$d(\gamma(t), \gamma_n(t)) \leq d(\gamma(t), \gamma'_n(t)) + d(\gamma_n(t), \gamma'_n(t)) \leq C(d(y, y_n) + d(x_n, x)) \rightarrow 0. \quad (3.1.7)$$

Since this estimate does not depend on t , we conclude that γ_n converges uniformly to γ .

(ii) Let $\gamma : [0, l] \rightarrow X$ be a local geodesic with $L(\gamma) \leq D_k$. Consider the set $T := \{t \in [0, l] \mid \gamma|_{[0, t]} \text{ is a geodesic}\}$. This set is closed and nonempty by definition. We show it to be open as well. Then connectedness of $[0, l]$ implies that $T = [0, l]$ and so γ is a geodesic. To this end fix $s \in T \cap (0, l)$. Since γ is a local geodesic, we find a subinterval $[a, b] \subseteq [0, l]$ containing s such that $\gamma|_{[a, b]}$ is a geodesic. Now consider the image $\gamma([0, b])$ (which we do not yet know to be a geodesic segment). This splits into the geodesic segments $[\gamma(0), \gamma(s)]$ and $[\gamma(s), \gamma(b)]$. Consider the triangle $\Delta([\gamma(0), \gamma(s)], [\gamma(s), \gamma(b)], [\gamma(0), \gamma(b)])$. Again, we emphasize that at this point we do not know that the geodesic segment $[\gamma(0), \gamma(b)]$ corresponds to the image $\gamma([0, b])$. However, the segments $[\gamma(0), \gamma(s)]$ and $[\gamma(s), \gamma(b)]$ have the segment $[\gamma(s), \gamma(s)]$ in common and this is already enough to show that the corresponding comparison triangle in M_k^2 must be degenerate: let x and y be points on the sides adjacent to $\gamma(s)$ and close to $\gamma(s)$. Say $x \in [\gamma(0), \gamma(s)]$, $y \in [\gamma(s), \gamma(b)]$ with $d(x, \gamma(s)) < d(\gamma(a), \gamma(s))$ and $d(y, \gamma(s)) < d(\gamma(a), \gamma(s))$. If it were not degenerate, then we would have by the CAT(k) inequality

$$d(x, y) \leq d(\bar{x}, \bar{y}) < d(\bar{x}, \overline{\gamma(s)}) + d(\overline{\gamma(s)}, \bar{y}) = d(x, \gamma(s)) + d(\gamma(s), y) = d(x, y), \quad (3.1.8)$$

where the last equality holds since x and y lie on the geodesic segment $[\gamma(a), \gamma(b)]$. This clearly cannot happen so the comparison triangle must be degenerate and consequently $\gamma([0, b]) = [\gamma(0), \gamma(b)]$. This implies that $b \in T$ and since T is obviously connected we have $(a, b) \subseteq T$ which shows that T is open, as desired.

(iii) Let $z \in X$, $0 < r < \frac{D_k}{2}$ and $x, y \in B_r(z) \subseteq X$. Consider the triangle $\Delta(x, y, z)$. Let $\bar{\Delta}(\bar{x}, \bar{y}, \bar{z})$ be its comparison triangle in M_k^2 . By construction we have $\bar{x}, \bar{y} \in B_r(\bar{z}) \subseteq M_k^2$. Since $B_r(\bar{z})$ is convex, we have $[\bar{x}, \bar{y}] \subseteq B_r(\bar{z})$. Now let $p \in [x, y]$ and $\bar{p} \in [\bar{x}, \bar{y}]$ be its comparison point. We then infer from the CAT(k) inequality that $d(z, p) \leq d(\bar{z}, \bar{p}) \leq r$. Thus, $p \in B_r(z)$ and consequently $[x, y] \subseteq B_r(z)$, showing that $B_r(z)$ is convex. The proof for the closed ball is completely analogous. $\square$

3.1.6 Remark (Diameter restriction important for continuous variation). The following example shows that the restriction on length is necessary for uniform convergence of geodesics. Consider M_1^2 , i.e. the “usual” 2-dimensional sphere and two opposite points, say x and $-x$. Clearly, there is no unique geodesic connecting these two points. Moreover, if we fix one great arc and consider sub arcs of the great arc opposite of the one we fixed, then these curves will not even converge pointwise (apart from the endpoints, obviously). More precisely, fix $y \in M_1^2$ unequal to x and consider the great arc from x to $-x$ in direction y . On the other side, we consider a sequence of great arcs from x to $-x$ in direction $-y$ but with its domain restricted to $[\frac{1}{n}, \pi - \frac{1}{n}]$. If we denote the endpoints of these curves by x_n and $-x_n$, then clearly $x_n \rightarrow x$ and $-x_n \rightarrow -x$. But (after we linearly parameterize these curves), the points in the interior of the (image of the) curves at corresponding times are “as far away as possible” from each other and in particular do not converge pointwise, not to mention uniformly.

Up until now we imagined a $\text{CAT}(k)$ space to have thinner triangles than a model space by requiring that points on the sides of any triangle are closer to each other than points in the comparison triangle. There are several other conditions one can impose on a geodesic space that essentially also amount to the space having thin triangles. For example, in a thin triangle, also the sum of angles should be smaller than the sum of angles of the corresponding comparison triangle. We will introduce different definitions and show that they are all equivalent to our space being $\text{CAT}(k)$. Afterwards, we will refer to any of these reformulations as the $\text{CAT}(k)$ inequality and use whichever one best fits the current situation. Note that the following characterizations all deal with triangles. With a bit more preparation afterwards, we will also develop a reformulation concerning the geometry of quadrilaterals.

3.1.7 Theorem (Equivalent conditions for being $\text{CAT}(k)$). *Let $k \in \mathbb{R}$ and let X be a D_k -geodesic space. We assume all triangles in this theorem to be of circumference $< 2D_k$. Then the following are equivalent:*

- (i) X is a $\text{CAT}(k)$ space.
- (ii) For every triangle $\Delta([x, y], [y, z], [z, x])$ in X and all $p \in [x, y]$ we have $d(z, p) \leq d(\bar{z}, \bar{p})$. This is similar to the formulation in 2.1.1.
- (iii) For every triangle $\Delta([x, y], [y, z], [z, x])$ in X and all $p \in [x, y], q \in [y, z]$ with $p \neq y \neq q$ we have $\angle_{\bar{y}}(\bar{p}, \bar{q}) \leq \angle_{\bar{y}}(\bar{x}, \bar{z})$. In other words, the vertex angle at $\bar{y}$ of the comparison triangle cannot increase when shortening the sides adjacent to y .
- (iv) For every triangle $\Delta([x, y], [y, z], [z, x])$ in X we have $\angle_x(y, z) \leq \angle_{\bar{x}}(\bar{y}, \bar{z}), \angle_y(x, z) \leq \angle_{\bar{y}}(\bar{x}, \bar{z})$ and $\angle_z(x, y) \leq \angle_{\bar{z}}(\bar{x}, \bar{y})$, i.e., we require that vertex angles in X are not greater than the corresponding comparison angles.
- (v) For every triangle $\Delta([x, y], [y, z], [z, x])$ in X with vertex angle γ at y we have: if $\bar{\Delta}(\bar{x}, \bar{y}, \bar{z})$ is a triangle in M_k^2 with $d(x, y) = d(\bar{x}, \bar{y}), d(y, z) = d(\bar{y}, \bar{z})$ and $\angle_{\bar{y}}(\bar{x}, \bar{z}) = \gamma$, then $d(x, z) \geq d(\bar{x}, \bar{z})$. This reminds us of the law of cosines: if we have a “comparison” triangle with equal vertex angle and adjacent sides, then the opposing side cannot be greater than in the original triangle.

Proof. (i) $\Rightarrow$ (ii) Clearly, (ii) is a special case of (i), so this follows immediately.

(ii) $\Rightarrow$ (iii) Let $\Delta([x, y], [y, z], [z, x])$ be a triangle in $X, p \in [x, y], q \in [y, z], p \neq y \neq q$. Let $T_1 := \bar{\Delta}(\bar{x}_1, \bar{y}_1, \bar{z}_1), T_2 := \bar{\Delta}(\bar{p}_2, \bar{y}_2, \bar{z}_2)$ and $T_3 := \bar{\Delta}(\bar{p}_3, \bar{y}_3, \bar{q}_3)$ be comparison triangles for $\Delta([x, y], [y, z], [z, x]), \Delta([p, y], [y, z], [z, p])$ and $\Delta([p, y], [y, q], [q, p])$, respectively. Let $\bar{q}_2 \in$

$[\bar{y}_2, \bar{z}_2] \subseteq T_2$ be the comparison point for q in T_2 . Then by (ii) we have $d(p, q) \leq d(\bar{p}_2, \bar{q}_2)$. By construction we have $d(\bar{p}_3, \bar{q}_3) = d(p, q)$ and so $d(\bar{p}_3, \bar{q}_3) \leq d(\bar{p}_2, \bar{q}_2)$. Consider now the triangle $T_4 := \bar{\Delta}(\bar{p}_2, \bar{y}_2, \bar{q}_2)$, i.e. T_2 with the side $[\bar{y}_2, \bar{z}_2]$ shortened to $[\bar{y}_2, \bar{q}_2]$. Two sides of this triangle are of the same length as their counterparts in T_3 . The third side in T_3 is shorter. Thus, the law of cosines gives that the angle at $\bar{y}_3$ in T_3 is smaller as well, i.e. $\angle_{\bar{y}_3}(\bar{p}_3, \bar{q}_3) \leq \angle_{\bar{y}_2}(\bar{p}_2, \bar{q}_2)$. Since $[\bar{y}_2, \bar{q}_2]$ is an initial segment of $[\bar{y}_2, \bar{z}_2]$ and the other side adjacent to $\bar{y}_2$ is equal in T_2 and T_4 , we conclude that by 1.2.9 the angle at $\bar{y}_2$ is the same in both triangles, i.e. $\angle_{\bar{y}_2}(\bar{p}_2, \bar{q}_2) = \angle_{\bar{y}_2}(\bar{p}_2, \bar{z}_2)$. In similar fashion, we again by (ii) have $d(p, z) \leq d(\bar{p}_1, \bar{z}_1)$ and by construction $d(p, z) = d(\bar{p}_2, \bar{z}_2)$. Now looking at the triangles $T_5 := \bar{\Delta}(\bar{p}_1, \bar{y}_1, \bar{z}_1)$ and T_2 we see that they share two sidelengths (the ones adjacent to $\bar{y}_1$ respectively $\bar{y}_2$) with the third side being shorter in T_2 . Thus, by the law of cosines the angle at $\bar{y}_2$ in T_2 is smaller than the angle at $\bar{y}_1$ in T_5 as well. Again, the construction of T_5 from T_1 did not change the angle at $\bar{y}_1$. So in the end we have shown $\angle_{\bar{y}_3}(\bar{p}_3, \bar{q}_3) \leq \angle_{\bar{y}_1}(\bar{x}_1, \bar{z}_1)$. Removing the indices, this is exactly the statement in (iii).

(iii) $\Rightarrow$ (iv) Recall that in 1.2.8 we defined the Alexandrov angle as a limit of vertex angles of smaller and smaller Euclidean triangles. Since by 2.3.1 the Alexandrov angle in M_k^2 agrees with the spherical respectively hyperbolic angle, we can replace the Euclidean triangles in the limiting process with corresponding triangles from an arbitrary model space without changing the value of the Alexandrov angle. With that in mind, (iv) easily follows from (iii): in a somewhat sloppy notation, we have $\angle_x(y, z) := \limsup_{p, q \rightarrow x} \angle_{\bar{x}}(\bar{p}, \bar{q})$, where p and q are points from $[x, y]$ and $[x, z]$, getting closer to x . Since by (iii) any such angle is bounded by $\angle_{\bar{y}}(\bar{x}, \bar{z})$, so is the lim sup, giving (iv).

(iv) $\Leftrightarrow$ (v) This is an immediate consequence of the law of cosines in M_k^2 : Let $\bar{\Delta}(\bar{x}, \bar{y}, \bar{z})$ be a comparison triangle in M_k^2 for $\Delta([x, y], [y, z], [z, x])$. By (iv) we have $\angle_y(x, z) \leq \angle_{\bar{y}}(\bar{x}, \bar{z})$. If we decrease the angle at $\bar{y}$, then the opposing side must become shorter by the law of cosines, giving (v). Conversely, if we have a triangle in M_k^2 as in (v), then shortening the side opposite of $\bar{y}$ causes the angle at $\bar{y}$ to decrease as well.

(iv) $\Rightarrow$ (ii) Let $\Delta([x, y], [y, z], [z, x])$ be a triangle in X . Let $p \in [x, y]$ and consider a geodesic segment $[p, z]$. Note that p splits $[x, y]$ into two segments, denoted by $[x, p]$ and $[p, y]$. Set $\gamma_1 := \angle_p([x, p], [p, z])$ and $\gamma_2 := \angle_p([p, z], [p, y])$. Set $\beta := \angle_y([x, y], [y, z])$ and let $\bar{\Delta}(\bar{x}_1, \bar{y}_1, \bar{z}_1)$ be a comparison triangle for $\Delta([x, y], [y, z], [z, x])$. Set $\beta_1 := \angle_{\bar{y}_1}(\bar{x}_1, \bar{z}_1)$. Now consider comparison triangles $\bar{\Delta}(\bar{x}_2, \bar{p}_2, \bar{z}_2)$ and $\bar{\Delta}(\bar{p}_2, \bar{y}_2, \bar{z}_2)$ for the triangles $\Delta([x, p], [p, z], [z, x])$ and $\Delta([p, y], [y, z], [z, p])$, respectively. Choose these comparison triangles such that they have the edge $[\bar{p}_2, \bar{z}_2]$ in common and such that the points $\bar{x}_2$ and $\bar{y}_2$ lie on opposite sides of the line extending this segment. Setting $\bar{\gamma}_1 := \angle_{\bar{p}_2}(\bar{x}_2, \bar{z}_2)$, $\bar{\gamma}_2 := \angle_{\bar{p}_2}(\bar{y}_2, \bar{z}_2)$ and $\bar{\beta}_2 := \angle_{\bar{y}_2}(\bar{x}_2, \bar{p}_2)$, we obtain by the triangle inequality for angles and the fact that the angle remains unchanged when restricting a geodesic segment

$$\gamma_1 + \gamma_2 \geq \angle_p([x, p], [p, y]) = \pi. \quad (3.1.9)$$

Thus, by (iv) we also have $\bar{\gamma}_1 + \bar{\gamma}_2 \geq \pi$. Applying Alexandrov's lemma 2.3.4 we obtain $\bar{\beta}_2 \leq \bar{\beta}_1$ and so with the law of cosines, $d(\bar{p}_1, \bar{z}_1) \geq d(\bar{p}_2, \bar{z}_2) = d(p, z)$ follows, giving (ii).

(iii) $\Rightarrow$ (i) Consider a triangle $\Delta([x, y], [y, z], [z, x])$ in X with p, q as in (iii). Let $\bar{\Delta}(\bar{x}_1, \bar{y}_1, \bar{z}_1)$ and $\bar{\Delta}(\bar{p}_2, \bar{y}_2, \bar{q}_2)$ be comparison triangles for $\Delta([x, y], [y, z], [z, x])$ and $\Delta([p, y], [y, q], [q, p])$, respectively. Let $\bar{p}_1 \in [\bar{x}_1, \bar{y}_1]$ and $\bar{q}_1 \in [\bar{y}_1, \bar{z}_1]$ be comparison points for p and q in $\bar{\Delta}(\bar{x}_1, \bar{y}_1, \bar{z}_1)$. Consider the vertex angles $\angle_{\bar{y}_1}(\bar{x}_1, \bar{z}_1)$ and $\angle_{\bar{y}_2}(\bar{p}_2, \bar{q}_2)$ at $\bar{y}_1$ and $\bar{y}_2$, respectively. By (iii) we have

$\bar{\angle}_{\bar{y}_1}(\bar{x}_1, \bar{z}_1) \geq \bar{\angle}_{\bar{y}_2}(\bar{p}_2, \bar{q}_2)$. Thus, the law of cosines implies $d(\bar{p}_1, \bar{q}_1) \geq d(\bar{p}_2, \bar{q}_2) = d(p, q)$, and we are done. $\square$

3.2 More comparisons and $\text{CAT}(k)$ hierarchy

As announced before, there is one more characterization we would like to cover. Our sources remain the same as in the last section.

3.2.1 Definition ($\text{CAT}(k)$ 4-point condition). *Let X be a length space, $k \in \mathbb{R}$. A quadruple x_1, y_1, x_2, y_2 of points in X is said to satisfy the $\text{CAT}(k)$ 4-point condition if the following holds: Consider the triangles $\Delta([x_1, y_1], [y_1, x_2], [x_2, x_1])$ and $\Delta([x_1, y_2], [y_2, x_2], [x_2, x_1])$ with common edge $[x_1, x_2]$ and their comparison triangles $\bar{\Delta}(\bar{x}_1, \bar{y}_1, \bar{x}_2)$ and $\bar{\Delta}(\bar{x}_1, \bar{y}_2, \bar{x}_2)$, respectively. Then for all $\bar{z} \in [\bar{x}_1, \bar{x}_2]$ it holds that*

$$d(\bar{y}_1, \bar{z}) + d(\bar{z}, \bar{y}_2) \geq d(y_1, y_2). \quad (3.2.1)$$

If one of the triangles $\Delta([x_1, y_1], [y_1, x_2], [x_2, x_1])$ or $\Delta([x_1, y_2], [y_2, x_2], [x_2, x_1])$ has circumference greater than $2D_k$ then we set the $\text{CAT}(k)$ 4-point condition to be trivially satisfied.

The intuition should be similar to the $\text{CAT}(k)$ inequality: one can think of the 4-point condition as requiring that y_1 and y_2 are at most as far apart as $\bar{y}_1$ and $\bar{y}_2$, thus forcing the sides $[x_1, y_1], [y_1, x_2], [x_2, y_2], [y_2, x_1]$ to bend even more inwards than their comparison segments in the appropriate model space.

3.2.2 Proposition (4-point condition is equivalent to being $\text{CAT}(k)$). *Let $k \in \mathbb{R}$ and let X be a D_k -geodesic space. Then the following are equivalent:*

- (i) X is a $\text{CAT}(k)$ space.
- (ii) Every quadruple of points x_1, y_1, x_2, y_2 in X satisfies the $\text{CAT}(k)$ 4-point condition.

Proof. (i) $\Rightarrow$ (ii) Let X be a $\text{CAT}(k)$ space. Let $x_1, y_1, x_2, y_2 \in X$ such that y_1 and y_2 lie on opposite sides of the line through x_1 and x_2 . Consider the triangles $\Delta(x_1, y_1, x_2)$ and $\Delta(x_1, y_2, x_2)$ with common edge $[x_1, x_2]$ and their comparison triangles $\bar{\Delta}(\bar{x}_1, \bar{y}_1, \bar{x}_2)$ and $\bar{\Delta}(\bar{x}_1, \bar{y}_2, \bar{x}_2)$, respectively. We suppose the circumference is small enough since otherwise there is nothing to prove. Let $z \in [x_1, x_2]$ and $\bar{z} \in [\bar{x}_1, \bar{x}_2]$ be its comparison point. By the $\text{CAT}(k)$ inequality we have $d(y_1, z) \leq d(\bar{y}_1, \bar{z})$ and $d(z, y_2) \leq d(\bar{z}, \bar{y}_2)$. So in particular we obtain

$$d(\bar{y}_1, \bar{z}) + d(\bar{z}, \bar{y}_2) \geq d(y_1, z) + d(z, y_2) \geq d(y_1, y_2). \quad (3.2.2)$$

(ii) $\Rightarrow$ (i) This direction of the proof is more or less trivial and is best summed up by the fact that any triangle can be viewed as a degenerate quadrilateral. More precisely, let $x_1, y_1, x_2 \in X$ and consider a triangle $\Delta([x_1, y_1], [y_1, x_2], [x_2, x_1])$. Let $y_2 \in [x_1, x_2]$ and consider comparison triangles $\bar{\Delta}(\bar{x}_1, \bar{y}_1, \bar{x}_2)$ and $\bar{\Delta}(\bar{x}_1, \bar{y}_2, \bar{x}_2)$, which do share the edge $[\bar{x}_1, \bar{x}_2]$ (in fact, the latter triangle only consists of this edge). Then for $\bar{z} = \bar{y}_2$ the $\text{CAT}(k)$ 4-point condition reduces to

$$d(\bar{y}_1, \bar{z}) = d(\bar{y}_1, \bar{z}) + d(\bar{z}, \bar{y}_2) \geq d(y_1, y_2) = d(y_1, z), \quad (3.2.3)$$

which already is the $\text{CAT}(k)$ inequality since y_2 was arbitrary. $\square$

Next, we will discuss properties and connections between $\text{CAT}(k)$ spaces for different $k \in \mathbb{R}$. On the one hand, it seems very intuitive that a space satisfying some curvature bound from above also satisfies a weaker bound from above. On the other hand, the converse of this statement is

also true in some sense. This equivalence, while seeming clunky and inconvenient, can be used to improve the curvature bound on some spaces.

Proving this requires quite a bit of technical finesse, so we will first show a lemma assisting us in the proof. Yet again, we will use the Riemannian structure in M_k^2 .

3.2.3 Lemma (Comparison of comparison triangles). *Let $k, k' \in \mathbb{R}, k < k'$. Let $A, B, C \in M_k^2, A', B', C' \in M_{k'}^2$. Consider triangles $T_1 := \Delta(A, B, C)$ and $T_2 := \Delta(A', B', C')$ with side-lengths a, b, c and a', b', c' respectively. More precisely, we have $d(A, C) = d(A', C') =: b, d(B, C) = d(B', C') =: a$, and $d(A, B) =: c$ and $d(A', B') =: c'$. Furthermore, assume that T_1 has circumference $< 2D_{k'}$ and that $0 < \bar{\angle}_C^{(k)}(A, B) = \bar{\angle}_{C'}^{(k')}(A', B') < \pi$. Then $c' < c$.*

Proof. Let (r, θ) be polar coordinates on $B_{D_{k'}}(C) \subseteq M_k^2$ and $B_{D_{k'}}(C') \subseteq M_{k'}^2$, around C and C' , respectively. Note that this means that if $k \leq 0$, then we introduce the polar coordinates on the whole space, centered at C . Let $h : B_{D_{k'}}(C') \rightarrow B_{D_{k'}}(C)$, $(r, \theta) \mapsto (r, \theta)$, i.e. the map that sends a point in $M_{k'}^2$ with polar coordinates (r, θ) to the corresponding point with the same coordinates in M_k^2 . In particular, $h(C') = C$ and since $T_{C'}h$ is an isometry, we may assume without loss of generality that $h(A') = A$ and $h(B') = B$. Consequently, we have $T_1 \subseteq h(B_{D_{k'}}(C')) \subseteq B_{D_{k'}}(C)$.

Our goal now is to show that $L(h \circ \gamma) > L(\gamma)$ for all non-radial paths γ with respect to the polar coordinates. To do this, we will show $\|T_p h(v)\| \geq \|v\|$ with equality if and only if v is a radial vector. Fix r and define the function

$$f_r : \left(-\infty, \left(\frac{\pi}{r}\right)^2\right) \rightarrow \mathbb{R}, \quad f_r(k) = \begin{cases} r, & \text{if } k = 0, \\ \frac{1}{\sqrt{k}} \sin(\sqrt{k}r), & \text{if } k > 0, \\ \frac{1}{\sqrt{-k}} \sinh(\sqrt{-k}r), & \text{if } k < 0. \end{cases}$$

With this in mind, recall the representation of the Riemannian metric on M_k^2 in polar coordinates:

$$ds^2 = dr^2 + f_r(k)^2 d\theta^2. \quad (3.2.4)$$

Clearly, f_r is continuous since each case is continuous and $\lim_{k \nearrow 0} \frac{1}{\sqrt{-k}} \sinh(\sqrt{-k}r) = r = \lim_{k \searrow 0} \frac{1}{\sqrt{k}} \sin(\sqrt{k}r)$ and strictly decreasing since again all cases are. Now let $v = v_1 \frac{\partial}{\partial r}(r, \theta) + v_2 \frac{\partial}{\partial \theta}(r, \theta) \in T_p M_{k'}^2, p \in B_{D_{k'}}(C')$. Then

$$\|v\|^2 = v_1^2 + f_r(k')^2 v_2^2 \leq v_1^2 + f_r(k)^2 v_2^2 = \|T_p h(v)\|^2, \quad (3.2.5)$$

with equality if and only if $v_2 = 0$, i.e. v being radial. So we have shown that $T_p h$ strictly increases the norm of non-radial vectors and consequently h strictly increases the length of non-radial paths. Finally, consider again the geodesic segment $[A, B]$ in T_1 of length c . By the above, this corresponds to the image under h of a path from A' to B' with length $< c$ (since the segment is not radial). Since $d(A', B') = c'$, we must have $c' < c$ as well. $\square$

3.2.4 Theorem (CAT(k) hierarchy). *Let X be length space, $k \in \mathbb{R}$. Then X is a CAT(k) space if and only if it is a CAT(k') space for all $k' > k$. In particular, X is D_k -geodesic if and only if it is $D_{k'}$ -geodesic for all $k' > k$.*

Proof. Let us first show the part about the geodesicity of X . Clearly, if X is D_k -geodesic then it is $D_{k'}$ -geodesic for all $k' > k$ since $D_{k'} \leq D_k$. Conversely, assume X to be $D_{k'}$ -geodesic for all $k' > k$. Let $x, y \in X$ with $d(x, y) < D_k$. Then there exists a $k_0 \in (k, \infty)$ such that

$d(x, y) < D_{k_0} \leq D_k$: the statement is only non-trivial when $k > 0$, but then the function $(k, \infty) \rightarrow (0, D_k)$, $l \mapsto D_l = \frac{\pi}{\sqrt{l}}$ is decreasing and continuous and takes on every value in its range. Since X is D_{k_0} -geodesic and x and y were arbitrary, it follows that X is D_k -geodesic as well.

Now suppose X is a $\text{CAT}(k)$ space. We already established that X is $D_{k'}$ -geodesic for all $k' > k$. Let $\Delta(x, y, z)$ be a triangle in X with circumference $< 2D_{k'}$ and consider its comparison triangles $T_k := \bar{\Delta}(\bar{x}_k, \bar{y}_k, \bar{z}_k)$ and $T_{k'} := \bar{\Delta}(\bar{x}_{k'}, \bar{y}_{k'}, \bar{z}_{k'})$ in M_k^2 and $M_{k'}^2$, respectively. Now recall 3.2.3 which can be summed up by the statement that if two triangles in M_k^2 and $M_{k'}^2$ have two sides of equal length as well as an equal angle adjacent to those sides then the third side is shorter in the triangle in $M_{k'}^2$. By applying the law of cosines, we get that when enlarging the third side in $T_{k'}$ - namely to the point where it matches the third side in T_k - then the angle which was equal before must increase. Thus, we obtain that the angles in T_2 are not smaller than the angles in T_1 . Since the angles in $\Delta(x, y, z)$ are not greater than in T_1 , they are also not greater than in T_2 . Since $\Delta(x, y, z)$ (apart from having small enough circumference) was arbitrary, we conclude that X satisfies the $\text{CAT}(k')$ inequality.

Conversely, suppose X is a $\text{CAT}(k')$ space for all $k' > k$. Let $T_1 := \Delta([x, y], [y, z], [z, y])$ be a geodesic triangle in X with circumference $< 2D_k$. Fix $k' > k$ such that the circumference is still smaller than $2D_{k'}$. Set $a := d(x, y), b := d(y, z), c := d(z, x)$ and $\gamma := \angle_x([x, y], [x, z])$. Suppose k is nonnegative and recall characterization 3.1.7(v). Let $\bar{\Delta}(\bar{x}, \bar{y}, \bar{z})$ be a triangle in $M_{k'}^2$ with $d(\bar{x}, \bar{y}) = a, d(\bar{x}, \bar{z}) = c$ and $\angle_{\bar{x}}(\bar{y}, \bar{z}) = \gamma$. Then the $\text{CAT}(k')$ inequality implies $b = d(y, z) \geq d(\bar{y}, \bar{z}) =: \bar{b}$. Since for any triangle in $M_{k'}^2$ the law of cosines is an equality, we infer the following inequality by substituting $\bar{b}$ with b :

$$\cos(b\sqrt{k'}) \leq \cos(a\sqrt{k'}) \cos(c\sqrt{k'}) + \sin(a\sqrt{k'}) \sin(c\sqrt{k'}) \cos(\gamma). \quad (3.2.6)$$

In particular, this inequality is equivalent with the triangle T_1 satisfying the $\text{CAT}(k')$ inequality. If k is positive, then letting k' go to k we obtain the same inequality with k' replaced by k . If $k = 0$, we have in the limit $b^2 \leq a^2 + c^2 - 2ac \cos(\gamma)$. In both cases, we conclude that the triangle satisfies the $\text{CAT}(k)$ inequality. Since the triangle was chosen arbitrarily, X is a $\text{CAT}(k)$ space. For negative k , the argument is entirely analogous. $\square$

3.3 Final preparations and Reshetnyak's gluing theorem

In this section we finally prove the gluing theorem. Concerning sources, we will primarily rely on [BH99, Ch. II.11] and [AKP19a, Ch. 8.9].

Recalling some terminology from 1.4, Reshetnyak's gluing theorem essentially states that the amalgamation of two $\text{CAT}(k)$ spaces is again a $\text{CAT}(k)$ space, provided that the isometric subspaces are convex and complete. A bit more work needs to be done in order to prove the theorem. To gain some basic intuition, we first show a ridiculously simplified version of the gluing theorem.

3.3.1 Definition (Metric bouquet). *Let $(X_i)_{i \in I}$ be a family of metric spaces, $a_i \in X_i$. The metric bouquet of the family $(X_i)_{i \in I}$ is the space obtained by identifying each a_i in the disjoint union $\bigsqcup_{i \in I} X_i$. Using 1.4.8, we see that the metric bouquet is a special case of the amalgamation by setting $A_i = \{a_i\}$ and taking as isometries the functions $\phi_i : \{a_i\} \rightarrow \{a\} =: A, a_i \mapsto a$, where we view A as a singleton metric space equipped with the discrete metric.*

3.3.2 Proposition ($\text{CAT}(k)$ bouquet). *Let $(X_i)_{i \in I}$ be a family of $\text{CAT}(k)$ spaces, $A_i := \{a_i\} \subseteq X_i$. Then $X := \bigsqcup_A X_i$ is a $\text{CAT}(k)$ space.*

Proof. First we note that A_i is trivially proper and that 1.4.9(iii) can easily be generalized for D_k -geodesicity. Thus, X is a D_k -geodesic space. Also, we want to remind ourselves that a geodesic between points from some X_i does not leave this (part of the) space.

Let $T_1 := \Delta([x, y], [y, z], [z, x])$ be a triangle in X with circumference $< 2D_k$ and set $a := [a_i]$. We easily see that there are only three cases to consider. First, if x, y and z belong to the same X_i then clearly the $\text{CAT}(k)$ inequality is satisfied since X_i is a $\text{CAT}(k)$ space. Next, suppose $x, y \in X_i$ and $z \in X_j$ for some $i, j \in I$. Then $[z, x]$ and $[z, y]$ share an initial segment since they must pass through a to go from X_j to X_i and so $\angle_z([z, x], [z, y]) = 0$. Concerning the other two angles, they remain unchanged when considering a triangle $T_2 := \Delta([x, y], [y, a], [a, x])$ instead of T_1 , i.e. $\angle_x([x, y], [x, a]) = \angle_x([x, y], [x, z])$ and $(\angle_y([y, x], [y, a]) = \angle_y([y, x], [y, z])$. Since $T_2 \subseteq X_i$, we know that it satisfies the $\text{CAT}(k)$ inequality. A comparison triangle $\bar{T}_1$ for T_1 and a comparison triangle $\bar{T}_2$ for T_2 have a side of equal length, namely $[\bar{x}, \bar{y}]$. But the other two sides are longer in $\bar{T}_1$, by $d(a, z)$ to be precise. So by the law of cosines, the angles at $\bar{x}$ and $\bar{y}$ in $\bar{T}_1$ are not smaller than in $\bar{T}_2$, so T_1 satisfies the $\text{CAT}(k)$ inequality as well. Finally, if all three vertices belong to different spaces, then all geodesics pass through the common point. In particular, we have at each vertex the situation from before at z , making $\Delta([x, y], [y, z], [z, x])$ into a degenerate triangle with all angles being equal to zero. Thus, we have seen that any triangle satisfies the $\text{CAT}(k)$ inequality and consequently X is a $\text{CAT}(k)$ space. $\square$

We want to take one more intermediate step before proving the gluing theorem in its full generality. The most difficult argument in the general case is to see that the amalgamation of two D_k -geodesic spaces X_1 and X_2 is again a D_k -geodesic space. If the spaces are proper, we already know this to be the case and we only need to deal with the $\text{CAT}(k)$ inequality. However, even this argument needs one more building block. The following gluing lemma for triangles intuitively states that dividing a given triangle into two subtriangles which satisfy the $\text{CAT}(k)$ inequality already implies that the original triangles satisfies the $\text{CAT}(k)$ inequality as well.

3.3.3 Lemma (Gluing lemma for triangles). *Let X be a D_k -geodesic space, $k \in \mathbb{R}$. Let $\Delta([x, y], [y, z], [z, x])$ be a triangle in X with circumference $< 2D_k$. Let $p \in [y, z]$ and consider a geodesic segment $[x, p]$ as well as the two resulting subtriangles $T_1 := \Delta([x, y], [y, p], [p, x])$ and $T_2 := \Delta([x, z], [z, p], [p, x])$. If the angles in T_1 and T_2 are at most as large as the corresponding angles in appropriate comparison triangles, then the same holds for the angles of a comparison triangle for $\Delta([x, y], [y, z], [z, x])$.*

Proof. Clearly, the assumption on the angles is precisely 3.1.7(iv). Let $\bar{\Delta}(\bar{x}, \bar{y}, \bar{p})$ and $\bar{\Delta}(\bar{x}, \bar{z}, \bar{p})$ be comparison triangles for T_1 and T_2 such that they have the edge $[\bar{x}, \bar{p}]$ in common and such that $\bar{y}$ and $\bar{z}$ lie on opposite sides of the line extending the segment $[\bar{x}, \bar{p}]$. By the triangle inequality for angles, we have

$$\angle_p([p, y], [p, x]) + \angle_p([p, x], [p, z]) \geq \angle_p([p, y], [p, z]) = \pi. \quad (3.3.1)$$

Since T_1 and T_2 satisfy the $\text{CAT}(k)$ inequality, we have

$$\bar{\angle}_{\bar{p}}(\bar{y}, \bar{x}) + \bar{\angle}_{\bar{p}}(\bar{x}, \bar{z}) \geq \angle_p([p, y], [p, x]) + \angle_p([p, x], [p, z]) \geq \pi \quad (3.3.2)$$

as well.

Now recall Alexandrov's lemma, for which this gluing lemma is an excellent application. First, it is clear that the two comparison triangles for T_1 and T_2 with common edge $[\bar{x}, \bar{p}]$ removed form a quadrilateral in M_k^2 . Moreover, by the above considerations we have that the interior


angle at $\bar{p}$ of the quadrilateral is at least π . Now we “straighten out” this quadrilateral, as in the second claim of 2.3.4, and immediately conclude that this results in a comparison triangle for our original triangle $\Delta([x, y], [y, z], [z, x])$! Keeping the notation somewhat similar, we denote by $\Delta(\bar{x}', \bar{y}', \bar{z}')$ the obtained comparison triangle for $\Delta([x, y], [y, z], [z, x])$. Then $\angle_{\bar{y}'}(\bar{x}', \bar{z}') \geq \angle_{\bar{y}}(\bar{x}, \bar{p}) \geq \angle_y(x, p) = \angle_y(x, z)$ by Alexandrov’s lemma. Analogously, we obtain the remaining estimates. $\square$

Note that in the following result, while we require A_1 to be complete as a metric space on its own, i.e. with the restricted metric from X_1 , convexity only makes sense when viewing it as a subset of X_1 . The same holds for A_2 .

3.3.4 Theorem (Reshetnyak). *Let X_1 and X_2 be proper $\text{CAT}(k)$ spaces, A_1 and A_2 complete and D_k -convex subspaces isometric to some metric space A . Then $X := X_1 \sqcup_A X_2$ is a $\text{CAT}(k)$ space.*

Proof. First, recall that any closed subset of a proper metric space is again proper. Thus, 1.4.9(iii) applies and we infer that X is D_k -geodesic. It remains to show that any triangle with small enough circumference satisfies the $\text{CAT}(k)$ inequality.

Let $\Delta([x, y], [y, z], [z, x])$ be a triangle in X with circumference $< 2D_k$. Clearly, if all three vertices belong to either X_1 or X_2 then the triangle satisfies the $\text{CAT}(k)$ inequality since X_1 and X_2 are $\text{CAT}(k)$ spaces. In addition, if say $x, y \in X_1$ and $z \in A_2$, then z can be identified with a point in A_1 and the same argument applies as in the case before. Thus, the only interesting case to consider is, up to symmetry of course, a triangle with $x, y \in X_1$ and $z \in X_2 \setminus A_2$. Let $p \in [x, z] \cap A_1$ and $q \in [y, z] \cap A_1$, i.e. p and q are points still in X_1 but “close to” where the segments jump into X_2 . In particular, $[x, p] \subseteq [x, z] \cap X_1$ and $[y, q] \subseteq [y, z] \cap X_1$, i.e. the initial segments up to p and q , respectively, are entirely contained in X_1 . Similarly, the other part of each segment is entirely contained in X_2 , so we have $[p, z] \subseteq [x, z] \cap X_2$ and $[q, z] \subseteq [y, z] \cap X_2$.



Now 3.3.3 comes into play. We split our triangle into three parts and apply the gluing lemma twice. First, consider the subtriangles $T_1 := \Delta([x, p], [p, q], [q, x])$ and $T_2 := \Delta([p, q], [q, z], [z, p])$. Clearly, we have $T_1 \subseteq X_1$ and $T_2 \subseteq X_2$ so both satisfy the $\text{CAT}(k)$ inequality. Moreover, they share the edge $[p, q]$. By the gluing lemma we now infer that the combined triangle $T_3 :=$

$\Delta([x, q], [q, z], [z, x])$ satisfies the $\text{CAT}(k)$ inequality as well (namely, its angles are at most as large as the angles of a comparison triangle in M_k^2). Next, we look at T_3 and $T_4 := \Delta([x, y], [y, q], [q, x])$. For T_3 we just showed that it satisfies the $\text{CAT}(k)$ inequality and $T_4 \subseteq X_1$, so again both parts satisfy the $\text{CAT}(k)$ inequality. Finally, we apply the gluing lemma once more to conclude that the combined triangle, which is precisely $\Delta([x, y], [y, z], [z, x])$, satisfies the $\text{CAT}(k)$ inequality. $\square$

3.3.5 Remark (On notation and assumptions). A careful reader surely has noticed that in the above proof we decided to not use the slightly more precise notation of equivalence classes for points in X . Since no computations regarding the metric were necessary, it should make no difference in readability. In the final version of the proof, we will return to the notation used throughout 1.4.

To prove the gluing theorem for arbitrary $\text{CAT}(k)$ spaces, we need two more lemmata. The first one intuitively says that an approximate midpoint between two points almost lies on the geodesic segment connecting them, provided that the segment is unique.

3.3.6 Lemma (Estimating distance from midpoint to segment). *Let X be a proper geodesic space and assume that the geodesic segment $[x, z]$ joining x and z in X is unique. Then for all $\varepsilon > 0$ there exists a $\delta > 0$ such that if $d(x, y) + d(y, z) < d(x, z) + \delta$, then $d(y, [x, z]) < \varepsilon$.*

Proof. Let $\varepsilon > 0$ and set $S_\varepsilon := \{y \in X \mid d(y, [x, z]) = \varepsilon\}$. If there is a point a distance more than ε from $[x, z]$, then S_ε is already non-empty: suppose $d(y, [x, z]) = \eta > \varepsilon$. Since X is proper and $[x, z]$ is closed, there is a point on the segment that achieves the infimum in the definition of $d(y, [x, z])$, i.e. there is a $p \in [x, z]$, such that $d(y, p) = \eta$. Consider the segment $[y, p]$, which exists since X is geodesic. By (L1) and (L3) in 1.2.1, there exists a $y_0 \in [y, p]$ which has the desired property. Thus, the only possibility for S_ε to be empty is if $d(y, [x, z]) < \varepsilon$ for all $y \in X$, in which case the assertion is trivially satisfied.

So assume $S_\varepsilon \neq \emptyset$. Clearly, S_ε is closed and bounded, hence compact since X is proper. Next, we define a continuous function $f : S_\varepsilon \rightarrow \mathbb{R}$ by $f(y) := d(x, y) + d(y, z) - d(x, z)$. Set $\delta := \min_{y \in S_\varepsilon} f(y)$, which exists since S_ε is compact. Since the segment $[x, z]$ is unique, we even have $\delta > 0$. Now let $q \in X$ be such that $d(q, [x, z]) = \eta \geq \varepsilon$. In particular, $d(q, x) \geq \varepsilon$, so by construction we find a $y \in S_\varepsilon$ lying on the segment $[q, x]$ such that $d(x, q) = d(x, y) + d(y, q)$. Thus, $d(x, q) + d(q, z) = d(x, y) + d(y, q) + d(q, z) \geq d(x, y) + d(y, z)$ and consequently $d(x, q) + d(q, z) \geq d(x, z) + f(y) \geq d(x, z) + \delta$. The implication we just showed is logically equivalent to our statement. $\square$

The next lemma follows from 3.3.6 and 2.3.2 and states that in any model space, approximate midpoints are not far away from midpoints.

3.3.7 Lemma (Approximate midpoints are close to midpoints). *Let $k \in \mathbb{R}, l < D_k, \varepsilon > 0$. Then there exists a positive constant δ depending on k, l and ε such that for all $x, y \in M_k^2$ with $d(x, y) < l$ the following holds: if we denote the midpoint of x and y by m and $m' \in M_k^2$ is a point such that both $d(x, m') \leq \frac{1}{2}d(x, y) + \delta$ and $d(m', y) \leq \frac{1}{2}d(x, y) + \delta$, then $d(m, m') < \varepsilon$.*

Proof. First recall that by 2.3.2, the action of $\text{Isom}(M_k^2)$ is transitive on equidistant pairs of points. Thus, we may assume without loss of generality that $[x, y]$ is the initial segment of a fixed segment $[x, z]$ of length l . As in 3.3.6, let η be such that $d(x, p) + d(p, z) < d(x, z) + \eta$ implies $d(p, [x, z]) < \frac{\varepsilon}{2}$. Note that by 2.2.10 the segment $[x, z]$ is unique since $d(x, z) = l < D_k$.

Now by the triangle inequality, we have $d(x, m') + d(m', z) \leq d(x, m') + d(m', y) + d(y, z)$, where the right hand side by assumption is bounded from above by $d(x, y) + d(y, z) + 2\delta =$

$d(x, z) + 2\delta$. Recall that this equality holds since y lies on the segment $[x, z]$. Thus, if $2\delta < \eta$ then $d(x, m') + d(m', z) \leq d(x, z) + 2\delta < d(x, z) + \eta$ and consequently $d(m', [x, z]) < \frac{\varepsilon}{2}$. In particular, we find $q \in [x, z]$ such that

$$d(m', q) < \frac{\varepsilon}{2} - \alpha \quad (3.3.3)$$

for some $\alpha > 0$. Then we have $d(x, q) \leq d(x, m') + d(m', q) \leq \frac{1}{2}d(x, y) + \delta + \frac{\varepsilon}{2}$ and we obtain a similar estimate with respect to y , i.e. $d(q, y) \leq \frac{1}{2}d(x, y) + \delta + \frac{\varepsilon}{2}$. Finally, we distinguish two cases: if $q \in [x, m] \subseteq [x, z]$, then $d(q, y) = d(m, q) + d(m, y) = d(m, q) + \frac{1}{2}d(x, y)$. And if $q \in [m, z] \subseteq [x, z]$, then $d(x, q) = d(x, m) + d(m, q) = \frac{1}{2}d(x, y) + d(m, q)$. In both cases, we arrive at

$$d(m, q) \leq \delta + \frac{\varepsilon}{2}. \quad (3.3.4)$$

Thus, by choosing $\delta < \min\{\alpha, \frac{\eta}{2}\}$ we obtain the desired estimate. Namely,

$$d(m, m') \leq d(m, q) + d(q, m') < \delta + \frac{\varepsilon}{2} + \frac{\varepsilon}{2} - \alpha < \varepsilon \quad (3.3.5)$$

by (3.3.3) and (3.3.4). $\square$


3.3.8 Theorem (Generalized gluing). *Let X_1 and X_2 be $\text{CAT}(k)$ spaces, A_1 and A_2 complete and D_k -convex subspaces isometric to some metric space A . Then $X := X_1 \sqcup_A X_2$ is a $\text{CAT}(k)$ space.*

Proof. Thanks to all our previous work, the only thing left to prove is that X actually is a D_k -geodesic space. After this, we simply repeat the arguments from 3.3.4. Thus, given two points $[x], [y] \in X$ with $\tilde{d}([x], [y]) < D_k$ we need to find a geodesic connecting them. Recall the proof of 1.4.9(iii), where we verified the existence of an $a \in A$ such that $\tilde{d}([x], [y]) = d_1(x, \phi_1(a)) + d_2(\phi_2(a), y)$ in order to verify that X is a geodesic space, where ϕ_1 and ϕ_2 are isometries from A to A_1 and A_2 , respectively. Since we already know that there are segments connecting x and $\phi_1(a)$ as well as $\phi_2(a)$ and y , we piece them together and project onto X to obtain a geodesic from x to y . We want to proceed similarly in this proof, i.e. find an $a \in A$ such that $\tilde{d}([x], [y]) = d_1(x, \phi_1(a)) + d_2(\phi_2(a), y)$. Since $\tilde{d}([x], [y]) = \inf_{a \in A} d_1(x, \phi_1(a)) + d_2(\phi_2(a), y)$ by definition, there exists a sequence $(a_n)_{n \in \mathbb{N}}$ in A such that $d_1(x, \phi_1(a_n)) + d_2(\phi_2(a_n), y) - \tilde{d}([x], [y]) < \frac{1}{n}$. The set $\{d_1(x, \phi_1(a_n)) \mid n \in \mathbb{N}\}$ is bounded and hence contained in a compact subset of $\mathbb{R}$. This means that we can extract a converging subsequence, and for the sake of readability we may assume without loss of generality that already the sequence $(d_1(x, \phi_1(a_n)))_{n \in \mathbb{N}}$ converges to some number $l_1 \in \mathbb{R}$. In particular, this implies $d_2(y, \phi_2(a_n)) \rightarrow l_2 := \tilde{d}([x], [y]) - l_1$. We show now that $(a_n)_{n \in \mathbb{N}}$ is a Cauchy sequence. Then by the completeness of A , $a := \lim_{n \rightarrow \infty} a_n$ exists and has the desired property, finishing the proof.


First, we may assume without loss of generality that $l_1 \leq \frac{\tilde{d}([x], [y])}{2} < \frac{D_k}{2}$ (otherwise just switch X_1 and X_2). Fix an $l \in (\tilde{d}([x], [y]), D_k) \subseteq \mathbb{R}$. Let $\varepsilon > 0$. Combining 3.3.6 and 3.3.7, let $\delta > 0$ be such that for all $p, q, r \in M_k^2$ with $d(p, r) < l$ and $d(p, q) + d(q, r) < d(p, r) + \delta$, we have $d(q, [p, r]) < \varepsilon$. Shrinking δ if necessary, we may assume without loss of generality that $\tilde{d}([x], [y]) + \delta < l$. Returning to the sequence $(a_n)_{n \in \mathbb{N}}$, let $N \in \mathbb{N}$ be such that for all $n > N$ we have $d_1(x, \phi_1(a_n)) < \frac{D_k}{2}$ and

$$|d_1(x, \phi_1(a_n)) - l_1| < \varepsilon \quad (3.3.6)$$

as well as $d_1(x, \phi_1(a_n)) + d_2(\phi_2(a_n), y) < \tilde{d}([x], [y]) + \delta$. Now let $m, n > N$, then $d_A(a_n, a_m) = d_1(\phi_1(a_n), \phi_1(a_m)) \leq d_1(\phi_1(a_n), x) + d_1(x, \phi_1(a_m)) < \frac{D_k}{2} + \frac{D_k}{2} = D_k$. Thus, the connecting segment $[\phi_1(a_n), \phi_1(a_m)]$ is unique by 3.1.5. Moreover, since $\phi_1(a_n), \phi_1(a_m) \in A_1$ and A_1 is D_k -convex, the segment is entirely contained in A_1 . Analogously, $[\phi_2(a_n), \phi_2(a_m)] \subseteq A_2$ is unique as well and so we obtain unique triangles $T_1 := \Delta(x, \phi_1(a_n), \phi_1(a_m)) \subseteq X_1$ and $T_2 := \Delta(y, \phi_2(a_n), \phi_2(a_m)) \subseteq X_2$.



Next, observe that we can bound the circumference of T_1 and T_2 by $2D_k$ so that we obtain comparison triangles in M_k^2 . Denote these comparison triangles by $\bar{\Delta}(\bar{x}, \bar{a}_n, \bar{a}_m)$ and $\bar{\Delta}(\bar{y}, \bar{a}_n, \bar{a}_m)$, respectively, and arrange them such that they share the edge $[\bar{a}_n, \bar{a}_m]$ and such that $\bar{x}$ and $\bar{y}$ lie on opposite sides of the line extending the common edge $[\bar{a}_n, \bar{a}_m]$. Furthermore, we suppose these comparison triangles to be non-degenerate since the degenerate case is trivial to prove. Now consider the unique segment $[\bar{x}, \bar{y}] \subseteq M_k^2$. We distinguish two cases: either the segments $[\bar{x}, \bar{y}]$ and $[\bar{a}_n, \bar{a}_m]$ have a unique point $\bar{p} \in M_k^2$ in common, or their intersection is empty.



Suppose first that $[\bar{x}, \bar{y}] \cap [\bar{a}_n, \bar{a}_m] = \{\bar{p}\}$. Let $p \in [\phi_1(a_n), \phi_1(a_m)]$ be such that $\bar{p}$ is a comparison point for p , i.e. $d_1(\phi_1(a_n), p) = d(\bar{a}_n, \bar{p})$. By the CAT(k) inequality in X_1 and X_2 , respectively, we have

$$\tilde{d}([x], [y]) \leq \tilde{d}([x], [p]) + \tilde{d}([p], [y]) \leq d(\bar{x}, \bar{p}) + d(\bar{p}, \bar{y}) = d(\bar{x}, \bar{y}), \quad (3.3.7)$$

where the last equality holds since $\bar{p}$ lies on the segment $[\bar{x}, \bar{y}]$. This implies

$$d(\bar{x}, \bar{a}_n) + d(\bar{a}_n, y) = d_1(x, \phi_1(a_n)) + d_2(\phi_2(a_n), y) \leq \tilde{d}([x], [y]) + \delta \leq d(\bar{x}, \bar{y}) + \delta. \quad (3.3.8)$$

Thus, by definition of δ we have $d(\bar{a}_n, [\bar{x}, \bar{y}]) < \varepsilon$ and $d(\bar{a}_m, [\bar{x}, \bar{y}]) < \varepsilon$. Let $\bar{p}_n \in [\bar{x}, \bar{y}]$ and $\bar{p}_m \in [\bar{x}, \bar{y}]$ be points such that $d(\bar{a}_n, \bar{p}_n) < \varepsilon$ and $d(\bar{a}_m, \bar{p}_m) < \varepsilon$, respectively. Since by (3.3.6) the values $d(\bar{x}, \bar{a}_n) = d_1(x, \phi_1(a_n))$ and $d(\bar{x}, \bar{a}_m) = d_1(x, \phi_1(a_m))$ both differ at most ε from l_1 , we conclude that their difference is at most 2ε , i.e.

$$|d(\bar{x}, \bar{a}_n) - d(\bar{x}, \bar{a}_m)| \leq 2\varepsilon. \quad (3.3.9)$$

Now suppose without loss of generality that on the segment $[\bar{x}, \bar{y}]$ the point $\bar{p}_m$ lies closer than $\bar{p}_n$ to $\bar{x}$, i.e. $[\bar{x}, \bar{p}_m] \subseteq [\bar{x}, \bar{p}_n]$ and so $d(\bar{p}_n, \bar{p}_m) = d(\bar{x}, \bar{p}_n) - d(\bar{x}, \bar{p}_m)$. By the triangle inequality, we infer

$$\begin{aligned} d(\bar{x}, \bar{a}_m) &\leq d(\bar{x}, \bar{p}_m) + d(\bar{p}_m, \bar{a}_m) && \Longleftrightarrow \\ d(\bar{x}, \bar{a}_m) - d(\bar{p}_m, \bar{a}_m) &\leq d(\bar{x}, \bar{p}_m) && \Longleftrightarrow \\ -d(\bar{x}, \bar{p}_m) &\leq -d(\bar{x}, \bar{a}_m) + d(\bar{p}_m, \bar{a}_m). \end{aligned}$$


From this we furthermore obtain

$$\begin{aligned} d(\bar{p}_n, \bar{p}_m) &= d(\bar{x}, \bar{p}_n) - d(\bar{x}, \bar{p}_m) \leq d(\bar{x}, \bar{a}_n) + d(\bar{a}_n, \bar{p}_n) - d(\bar{x}, \bar{p}_m) \\ &\leq d(\bar{x}, \bar{a}_n) + d(\bar{a}_n, \bar{p}_n) - d(\bar{x}, \bar{a}_m) + d(\bar{p}_m, \bar{a}_m) \\ &= (d(\bar{x}, \bar{a}_n) - d(\bar{x}, \bar{a}_m)) + d(\bar{a}_n, \bar{p}_n) + d(\bar{a}_m, \bar{p}_m) \\ &< 2\varepsilon + \varepsilon + \varepsilon = 4\varepsilon. \end{aligned}$$

Consequently we arrive at

$$\begin{aligned} d_A(a_n, a_m) &= d_1(\phi_1(a_n), \phi_1(a_m)) = d(\bar{a}_n, \bar{a}_m) \\ &\leq d(\bar{a}_n, \bar{p}_n) + d(\bar{p}_n, \bar{p}_m) + d(\bar{p}_m, \bar{a}_m) < \varepsilon + 4\varepsilon + \varepsilon = 6\varepsilon, \end{aligned}$$

which shows that $(a_n)_{n \in \mathbb{N}}$ is a Cauchy sequence.



Finally, suppose $[\bar{x}, \bar{y}] \cap [\bar{a}_n, \bar{a}_m] = \emptyset$. In this case we will again make use of Alexandrov's lemma. View the comparison triangles for T_1 and T_2 as a quadrilateral by removing the common edge $[\bar{a}_n, \bar{a}_m]$. Since $[\bar{x}, \bar{y}]$ and $[\bar{a}_n, \bar{a}_m]$ do not intersect, the quadrilateral must be concave. In particular, we have $\angle_{\bar{a}_n}(\bar{x}, \bar{a}_m) + \angle_{\bar{a}_n}(\bar{a}_m, \bar{y}) \geq \pi$ or $\angle_{\bar{a}_m}(\bar{x}, \bar{a}_n) + \angle_{\bar{a}_m}(\bar{a}_n, \bar{y}) \geq \pi$, i.e. the vertex angle in the quadrilateral is at least π . Suppose without loss of generality this is the case for $\bar{a}_n$. By 2.3.4, via "straightening out" the quadrilateral we obtain a triangle $\bar{\Delta}(\bar{x}', \bar{y}', \bar{a}'_m)$ such that $d(\bar{x}', \bar{a}'_m) = d(\bar{x}, \bar{a}_m)$, $d(\bar{y}', \bar{a}'_m) = d(\bar{y}, \bar{a}_m)$ and $d(\bar{x}', \bar{y}') = d(\bar{x}, \bar{a}_n) + d(\bar{a}_n, \bar{y})$. Let $\bar{a}'_n \in [\bar{x}', \bar{y}']$ such that $d(\bar{x}', \bar{a}'_n) = d(\bar{x}, \bar{a}_n)$. Then 2.3.4(ii) gives

$$d(\bar{a}'_n, \bar{a}'_m) \geq d(\bar{a}_n, \bar{a}_m) = d_1(\phi_1(a_n), \phi_1(a_m)). \quad (3.3.10)$$

By definition of δ , we have $d(\bar{a}'_m, [\bar{x}', \bar{y}']) < \varepsilon$. Let $v \in [\bar{x}', \bar{y}']$ achieve this infimum, so $d(v, \bar{a}'_m) < \varepsilon$. By (3.3.9), we have $|d(\bar{x}', \bar{a}'_n) - d(\bar{x}', \bar{a}'_m)| = |d(\bar{x}, \bar{a}_n) - d(\bar{x}, \bar{a}_m)| < 2\varepsilon$. At last, we distinguish on the segment $[\bar{x}', \bar{y}']$ whether v is closer than $\bar{a}'_n$ to $\bar{x}'$ or vice versa. In particular, we either have $d(v, \bar{a}'_n) = d(\bar{x}', v) - d(\bar{x}', \bar{a}'_n)$ if $\bar{a}'_n$ is closer to $\bar{x}'$ or $d(v, \bar{a}'_n) = d(\bar{x}', \bar{a}'_n) - d(\bar{x}', v)$ if v is closer to $\bar{x}'$. On the one hand, if $d(v, \bar{a}'_n) = d(\bar{x}', v) - d(\bar{x}', \bar{a}'_n)$ we compute

$$\begin{aligned} d(\bar{a}'_n, \bar{a}'_m) &\leq d(\bar{a}'_n, v) + d(v, \bar{a}'_m) = d(\bar{x}', v) - d(\bar{x}', \bar{a}'_n) + d(v, \bar{a}'_m) \\ &\leq d(\bar{x}', \bar{a}'_m) + d(\bar{a}'_m, v) - d(\bar{x}', \bar{a}'_n) + d(v, \bar{a}'_m) \\ &= (d(\bar{x}', \bar{a}'_m) - d(\bar{x}', \bar{a}'_n)) + 2d(\bar{a}'_m, v) < 2\varepsilon + 2\varepsilon = 4\varepsilon. \end{aligned}$$

On the other hand, if $d(v, \bar{a}'_n) = d(\bar{x}', \bar{a}'_n) - d(\bar{x}', v)$, first note that $d(\bar{x}', \bar{a}'_n) \leq d(\bar{x}', \bar{a}'_m) + 2\varepsilon$ by (3.3.9), as well as

$$d(\bar{x}', \bar{a}'_m) \leq d(\bar{x}', v) + d(v, \bar{a}'_m) \iff d(\bar{x}', \bar{a}'_m) - d(\bar{x}', v) \leq d(v, \bar{a}'_m). \quad (3.3.11)$$

Then we compute

$$\begin{aligned} d(\bar{a}'_n, \bar{a}'_m) &\leq d(\bar{a}'_n, v) + d(v, \bar{a}'_m) = d(\bar{x}', \bar{a}'_n) - d(\bar{x}', v) + d(v, \bar{a}'_m) \\ &\leq d(\bar{x}', \bar{a}'_m) - d(\bar{x}', v) + d(v, \bar{a}'_m) + 2\varepsilon \\ &\leq d(\bar{a}'_m, v) + d(v, \bar{a}'_m) + 2\varepsilon < \varepsilon + \varepsilon + 2\varepsilon = 4\varepsilon, \end{aligned}$$

which together with (3.3.10) shows that $(a_n)_{n \in \mathbb{N}}$ is a Cauchy sequence, and we are done. $\square$

As a last result, we want to generalize the gluing theorem to arbitrary families of $\text{CAT}(k)$ spaces. Having the previous theorem at hand, there is very little to do.

3.3.9 Theorem (Gluing arbitrary $\text{CAT}(k)$ families). *Let $(X_i)_{i \in I}$ be a family of $\text{CAT}(k)$ spaces, A a complete $\text{CAT}(k)$ space and $A_i \subseteq X_i$ closed subspaces isometric to A . Note that in this case each A_i is complete as well as D_k -convex. Then $X := \bigsqcup_A X_i$ is a $\text{CAT}(k)$ space.*

Proof. First, we note that the inclusion $X_j \bigsqcup_A X_k \hookrightarrow X$ is an isometry for all $j, k \in I$. To see this, let $x, y \in X_j \cup X_k$ and recall (1.4.2). Denote the metrics in $X_j \bigsqcup_A X_k$ and X by $\tilde{d}_{j,k}$ and $\tilde{d}_X$, respectively. Observe that if both x and y belong to, say, X_j , then we have $\tilde{d}_X([x], [y]) = \tilde{d}_{j,k}([x], [y]) = d_j(x, y)$, so suppose $x \in X_j$ and $y \in X_k$. On the one hand, we immediately get $\tilde{d}_X([x], [y]) \leq \tilde{d}_{j,k}([x], [y])$, because there is a larger set of chains in X over which we take the infimum. On the other hand, as in 1.4.9 we see that any chain in X can be shortened by removing all points that do not belong to X_j or X_k , so we get $\tilde{d}_X([x], [y]) = \tilde{d}_{j,k}([x], [y])$, which shows that the inclusion is an isometry. Thus, if $\tilde{d}_X([x], [y]) < D_k$ we can simply take the

connecting geodesic in $X_j \sqcup_A X_k$ and project further onto X , which shows that X is D_k -geodesic as well.

Now let $\Delta([x, y], [y, z], [z, x])$ be a triangle in X . The only case not covered by 3.3.4 is if all three vertices belong to different spaces, say $x \in X_i, y \in X_j$ and $z \in X_k$. But by the above considerations, this triangle is contained in $(X_i \sqcup_A X_j) \sqcup_A X_k$, which is a $\text{CAT}(k)$ space by applying 3.3.8 twice. Hence, $\Delta([x, y], [y, z], [z, x])$ satisfies the $\text{CAT}(k)$ inequality and the proof is finished. $\square$

3.3.10 Remark (Gluing spaces with curvature bounded below). As a closing remark to this chapter, we want to briefly talk about the case of curvature bounds from below in connection with gluing constructions. As was hinted at in 2.1.1(ii), we can turn the tables completely and basically reverse all inequalities mentioned in this chapter. If these reversed inequalities hold for all triangles in a D_k -geodesic space X and corresponding comparison triangles in M_k^2 , we say that X has curvature bounded below by k , or that X is a $\text{CBB}(k)$ space. While there certainly are nice constructions and theorems about spaces of this kind, there is no hope to get a similar result to the gluing theorem of Reshetnyak. Indeed, already the metric bouquet of a family of $\text{CBB}(k)$ spaces is in general not a $\text{CBB}(k)$ space. In 3.3.2, we saw that any triangle which has vertices in different spaces has vanishing angles, since the initial geodesic segments coincide, while the comparison triangle need not be degenerate. Thus, the angles of such a triangle cannot be larger than the angles of a comparison triangle.

4 Puff pastry and billiards

This last chapter is largely motivated by [AKP19b, Ch. 2] and gives an outlook on the somewhat astonishing applications of the gluing theorem. In addition we will make use of [BFK98b] and [BFK98a] as well as [Val64] and [Roc70]. Our goal in this chapter is to prove the collision theorem, which gives an upper bound on the number of collisions of a billiard trajectory with a billiard table, which is defined as the complement of a union of convex subsets (with nonempty intersection) in $\mathbb{R}^n$. From this we will derive a corollary which bounds the number of collisions of several balls. The setting of the corollary actually somewhat resembles a situation one might encounter when playing a real game of billiards.

4.1 Preliminaries

We start this section by defining a construction which will be of great importance throughout the whole chapter. Informally, we consider a (finite) sequence of copies of $\mathbb{R}^n$ and glue successive pairs of spaces along closed, convex subsets. Let $m, n \in \mathbb{N}$ and consider a sequence of $m+1$ copies of $\mathbb{R}^n$. We denote the i -th copy of $\mathbb{R}^n$ in this sequence by R_i . While this seems unnecessary now it will drastically simplify notation later on. Let $A^1, \dots, A^m$ be closed, convex subsets in $\mathbb{R}^n$.

As a first step, construct $\mathcal{R}_1 := R_1 \sqcup_{A_1} R_2$, i.e. two copies of $\mathbb{R}^n$ glued along the subset A_1 . Clearly, this process is the amalgamation of R_1 and R_2 , where the identified sets are not only isometric but equal (simply set $A_2 = A_1, A = A_1$ and $I = \{1, 2\}$ in 1.4.8). Next, we want to glue $\mathcal{R}_1$ and R_3 along A_2 , but in such a way that A_2 in $\mathcal{R}_1$ is only considered as a subset of R_2 . To put it in another way, the only points from $\mathcal{R}_1$ identified with points from R_3 should have originally belonged to R_2 . We denote this construction by $\mathcal{R}_2$. Formally, this could be achieved by setting $A^2 := \{[(a, 2)] \mid a \in A_2\}$, i.e. A^2 is the set of equivalence classes of points in A_2 viewed as a subset of R_2 in the disjoint union $R_1 \sqcup R_2$. Clearly, A^2 and A_2 are isometric, so by setting

$X_1 = \mathcal{R}_1, X_2 = \mathcal{R}_3, A_1 = A^2, A_2 = A_2$ and $A = A_2$, we can also write $\mathcal{R}_2$ as an amalgamation, namely $\mathcal{R}_2 := \mathcal{R}_1 \sqcup_{A_2} \mathcal{R}_3 = (R_1 \sqcup_{A_1} R_2) \sqcup_{A_2} R_3$. Now we continue the construction recursively. So we obtain $\mathcal{R}_i := \mathcal{R}_{i-1} \sqcup_{A_i} R_{i+1}$, where we consider A_i in $\mathcal{R}_i$ only as a subset of R_i , realized as $A^i := \{[(a, i)] \mid a \in A_i\}$. The final space $\mathcal{R}_m$ in this process will be denoted by $\mathcal{P}$.


4.1.1 Definition (Reshetnyak's puff pastry). *Let $m, n \in \mathbb{N}$ and suppose we are given a construction as described above. We combine the subsets $A_1, \dots, A_m$ into a sequence $\mathcal{A} := (A_1, \dots, A_m)$.*

- (i) *We call $\mathcal{A}$ an array and say that $\mathcal{P}$ is a Reshetnyak's puff pastry (or just puff pastry) for the array $\mathcal{A}$.*
- (ii) *The copies of $\mathbb{R}^n$ in the puff pastry will be called levels. It will not be necessary to distinguish between the i -th copy of $\mathbb{R}^n$ and its projection into $\mathcal{P}$, so we will denote both by R_i . If $x \in \mathbb{R}^n$, then we will denote by x^i the corresponding point in R_i .*
- (iii) *A puff pastry $\mathcal{P}$ is called end-to-end convex, if $R_1 \cup R_{m+1}$ is a convex subset in $\mathcal{P}$.*

4.1.2 Remark (On notation). Let $\mathcal{P}$ be a puff pastry for an array $\mathcal{A} = (A_1, \dots, A_m), x, y \in \mathbb{R}^n$ and assume without loss of generality that $i \leq j$. Then we have $x^i = y^j$ if and only if $x = y$ and either $i = j$ or $x \in \bigcap_{k=i}^{j-1} A_k$. This allows us to denote an arbitrary point in $\mathcal{P}$ by x^i and omit the use of equivalence classes. Additionally, when talking about different segments in several puff pastries, if confusion is imminent, we may emphasize to which puff pastry a segment belongs. For example, if $\mathcal{P}$ and $\mathcal{P}'$ are two puff pastries, we denote geodesic segments by $[x^i, y^j]_{\mathcal{P}}$ and $[x^{i'}, y^{j'}]_{\mathcal{P}'}$, respectively.

To better digest this somewhat dry and technical construction, we give an easy example in 1 dimension.

4.1.3 Example (Puff pastry for beginners). We consider four copies of $\mathbb{R}$ which we will denote by R_1, R_2, R_3 and R_4 , respectively. Let $A_1 = [1, 2], A_2 = [3, 4]$ and $A_3 = [0, 1]$. Then the puff pastry $\mathcal{P}$ corresponding to $\mathcal{A} = (A_1, A_2, A_3)$ looks like the following diagram.



As is easily seen, the name of this construction is definitively justified. Also, we see that this puff pastry is not end-to-end convex. Let $x = -1$, then any path connecting x^1 and x^4 passes through R_2 and R_3 , so $R_1 \cup R_4$ is not a convex subset of the puff pastry. In fact, $\bigcap_{i=1}^m A_i \neq \emptyset$ is a necessary condition on $\mathcal{A}$ for $\mathcal{P}$ to be end-to-end convex.

Now we reap the first rewards of the gluing theorem and our tedious notation in the construction of puff pastries above. Namely, one easily observes that any puff pastry is a $\text{CAT}(0)$ space.

4.1.4 Corollary (Puff pastry is CAT(0)). *Let $m, n \in \mathbb{N}$ and let $\mathcal{P}$ be a puff pastry for an array $\mathcal{A} = (A_1, \dots, A_m)$ of closed, convex subsets of $\mathbb{R}^n$. Then $\mathcal{P}$ is a CAT(0) space.*

Proof. Clearly, $\mathbb{R}^n$ is a CAT(0) space, and each A_i is closed and convex by assumption. Thus, we simply apply 3.3.4 m times, which we can do because in the above description we have seen that the puff pastry is constructed from several amalgamations. $\square$

Next, we want to show some properties of puff pastries which are similar in the sense that the sets in the first array are contained in the sets of the second array.

4.1.5 Proposition (Properties of similar puff pastries). *Let $\mathcal{A} = (A_1, \dots, A_m)$ and $\mathcal{A}' = (A'_1, \dots, A'_m)$ be two arrays of closed, convex subsets, such that $A_i \subseteq A'_i$ for all $i = 1, \dots, m$. We will write $\mathcal{A} \subseteq \mathcal{A}'$ for such a constellation. Let $\mathcal{P}$ and $\mathcal{P}'$ be the corresponding puff pastries.*

- (i) *The map $\phi : \mathcal{P} \rightarrow \mathcal{P}'$ defined by $x^i \mapsto x^{i'}$ is nonincreasing, i.e. we have $d_{\mathcal{P}}(x^i, y^j) \geq d_{\mathcal{P}'}(x^{i'}, y^{j'})$ for all $x^i, y^j \in \mathcal{R}$.*
- (ii) *Moreover, if for some $x^i, y^j \in \mathcal{R}$ the map ϕ preserves the distance, i.e. $d_{\mathcal{R}}(x^i, y^j) = d_{\mathcal{R}'}(x^{i'}, y^{j'})$, then the geodesic segments also correspond to each other. That is, the segment $[x^{i'}, y^{j'}]$ is the image of $[x^i, y^j]$ under ϕ .*

Proof. (i) This follows directly from the definition of the metric of the amalgamation, cf. (1.4.8): the sets with respect to which we identify points in $\mathcal{P}$ are contained in the corresponding sets for $\mathcal{P}'$. Thus, the set over which we take the infimum for $d_{\mathcal{P}}(x^i, y^j)$ is contained in the corresponding set for $d_{\mathcal{P}'}(x^{i'}, y^{j'})$, so we get $d_{\mathcal{P}}(x^i, y^j) \geq d_{\mathcal{P}'}(x^{i'}, y^{j'})$.

(ii) By 4.1.4 and 3.1.5, $\mathcal{P}$ and $\mathcal{P}'$ are uniquely geodesic spaces. Since we have $d_{\mathcal{R}}(x^i, y^j) = d_{\mathcal{R}'}(x^{i'}, y^{j'})$, the lengths of the segments $[x^{i'}, y^{j'}]$ and $[x^i, y^j]$ are equal, and because the segments are unique and ϕ is nonincreasing by (i), the second statement follows. $\square$

4.2 Criteria for end-to-end convexity

Next, we take a quick look at some criteria for end-to-end convexity of puff pastries. The first one implies the end-to-end convexity of a puff pastry that is constructed in a certain way from another puff pastry, provided that the first puff pastry is end-to-end convex.

4.2.1 Proposition (Criterion for end-to-end convexity I). *Let $\mathcal{A} = (A_1, \dots, A_m)$ be an array of closed and convex subsets of $\mathbb{R}^n$. Let $\mathcal{A}'$ be an array that is constructed by adding to $\mathcal{A}$ several subsets which already appeared in $\mathcal{A}$. More precisely, we have $\mathcal{A}' = (A'_1, \dots, A'_M)$, where $A'_i \in \{A_1, \dots, A_m\}$ for all $i = 1, \dots, M$ and there is a subset $\{i_1, \dots, i_m\} \subseteq \{1, \dots, M\}$, $i_1 < i_2 < \dots < i_m$, such that $A'_{i_1} = A_1, A'_{i_2} = A_2, \dots, A'_{i_m} = A_m$. This last requirement ensures that the original order of $\mathcal{A}$ is preserved in $\mathcal{A}'$. For example, if $\mathcal{A} = (A_1, A_2, A_3)$, then a possibility for $\mathcal{A}'$ would be $(A_1, A_2, A_2, A_1, A_3, A_2, A_3)$, but an illegal choice would be $(A_2, A_3, A_3, A_2, A_1)$. Denote by $\mathcal{P}$ and $\mathcal{P}'$ the corresponding puff pastries. If $\mathcal{P}$ is end-to-end convex, then $\mathcal{P}'$ is end-to-end convex as well.*

Proof. First, note that it suffices to show the statement for an array $\mathcal{A}'$ where we only add one extra subset from $\mathcal{A}$. Once we have proven this, we can use it to inductively obtain the general result. More precisely, suppose $\mathcal{P}$ is end-to-end convex and $\mathcal{P}_1$ is a puff pastry corresponding to an array $\mathcal{A}_1$, which has one more entry than $\mathcal{A}$ and we already showed that $\mathcal{P}_1$ is end-to-end convex. Then we use $\mathcal{P}_1$ to show the statement for some $\mathcal{P}_2$ which has one more

level than $\mathcal{P}_1$. After a finite amount of steps, we have constructed $\mathcal{P}'$ in this way and are finished.

So suppose without loss of generality that $\mathcal{A}' = (A_1, \dots, A_{i-1}, B, A_i, \dots, A_m)$, where $B \in \{A_1, \dots, A_m\}$. Let $\mathcal{A}'' = (A_1, \dots, A_{i-1}, \mathbb{R}^n, A_i, \dots, A_m)$ and let $\mathcal{P}''$ be the corresponding puff pastry. Then $\mathcal{P}''$ and $\mathcal{P}$ are isometric and we have $\mathcal{A}' \subseteq \mathcal{A}''$, so we can apply 4.1.5 to $\mathcal{P}'$ and $\mathcal{P}''$. Let $x^{1'}$ and $y^{m+2'}$ be points in $\mathcal{P}'$ in the first and last level, respectively. Let x^1 and y^{m+1} be the corresponding points in $\mathcal{P}$. Note that $\mathcal{P}$ has one fewer level than $\mathcal{P}'$, so y^{m+1} is in the last level in $\mathcal{P}$. Once the segment $[x^1, y^{m+1}]$ enters from R_1 into $\bigcap_{i=1}^m A_i$, it stays there until leaving into R_{m+1} , because $\mathcal{P}$ is end-to-end convex and $\bigcap_{i=1}^m A_i$ is convex as the intersection of convex sets. Since $\bigcap_{i=1}^m A_i \cap B = \bigcap_{i=1}^m A_i$, the same holds in $\mathcal{P}'$, so we have that the minimising geodesic corresponding to $[x^1, y^{m+1}]$ in $\mathcal{P}$ is also a valid connecting curve in $\mathcal{P}'$. This means that the distance of the endpoints in $\mathcal{P}'$ is at most as large as the distance in $\mathcal{P}$, so we have $d_{\mathcal{P}}(x^1, y^{m+1}) \geq d_{\mathcal{P}'}(x^{1'}, y^{m+2'})$. But by 4.1.5(i), we have $d_{\mathcal{P}}(x^1, y^{m+1}) = d_{\mathcal{P}''}(x^{1''}, y^{m+2''}) \leq d_{\mathcal{P}'}(x^{1'}, y^{m+2'})$, and so the distances are equal. But then by 4.1.5(ii), the segments correspond to each other, which ensures that $[x^{1'}, y^{m+2'}]$ stays in $R'_1 \cup R'_{m+2} \subseteq \mathcal{P}'$, and so $\mathcal{P}'$ is end-to-end convex as well. $\square$


We want to show another criterion for end-to-end convexity using dihedral angles. Recall from elementary geometry that any subset of $\mathbb{R}^n$ of the form $\{x \in \mathbb{R}^n \mid \langle x, v \rangle \geq b\}$, where $b \in \mathbb{R}$ and $v \in \mathbb{R}^n$, is called a (closed) half-space. Reversing the inequality one obtains again a half-space and their intersection is a hyperplane.

4.2.2 Definition (Dihedral angle). *Let $n \in \mathbb{N}$ and consider two half-spaces H_1 and H_2 in $\mathbb{R}^n$.*

- (i) *The intersection of two half-spaces $D := H_1 \cap H_2 \subseteq \mathbb{R}^n$ is called a dihedral angle. The intersection of the corresponding hyperplanes is called the edge of D . To this edge one can assign a value similar to the usual meaning of an angle. Namely, if v_1 and v_2 are the normal vectors corresponding to H_1 and H_2 , respectively, then the angle measure of D is the number $\alpha \in [0, \pi)$, such that*


$$\cos(\alpha) = \frac{\langle v_1, v_2 \rangle}{\|v_1\| \|v_2\|}. \quad (4.2.1)$$

- (ii) *Let $X \subseteq \mathbb{R}^n$ be a convex subset and let $p \in X$. A dihedral angle D is said to support X at p if $X \subseteq D$ and p belongs to the edge of D .*



The second criterion for end-to-end convexity is based on a lemma, which we will motivate almost playfully by an example.

4.2.3 Example (Elevator vs. staircase). Consider a situation as described in the diagram below. Suppose one wants to travel from point p to point q via one of the two drawn paths.



It is fairly obvious that the path through the intersection of the thick lines is longer for any p, q in the same connected component of the plane. Now we impose the additional constraint that on the shorter path one has to traverse the middle piece of the path several times. With enough fantasy, one can translate this situation into that of a building where one has two options going from p on the top floor to q on the bottom floor: going via the staircase represented by the middle part of the at first shorter path, or walking on the top floor from p to a modern, instantly teleporting elevator and then on the bottom floor to q . If one descends only a few floors, the stairs will be faster than the elevator. But clearly, if the building is tall enough and has enough floors, it will be faster to take the elevator. As it turns out, the following lemma states that there is a number of floors depending only on the angle α such that no matter where p and q are (on the top and bottom floor, respectively), it will always be faster to take the elevator. This building can be described more rigorously by a puff pastry for an array consisting of two alternating half-spaces, and the fact that the elevator is faster simply means that the puff pastry is end-to-end convex.

The proof of the lemma is greatly simplified by thinking of the puff pastry in a slightly different way, for which we need one more definition.

4.2.4 Definition (Metric cone). *Let (X, d) be a metric space. The metric cone of X is the space $\text{Con}(X) := (X \times [0, \infty)) / \sim$, where $(x, 0) \sim (y, 0)$ for all $x, y \in X$. This point is called the tip or the apex of the cone and is denoted by 0_X . The metric, defined via*

$$d_c((x, t), (y, s)) := \sqrt{t^2 + s^2 - 2ts \cos(\alpha)}, \quad (4.2.2)$$

where $\alpha = \min\{d(x, y), \pi\}$, is called the cone metric on $\text{Con}(X)$.


One motivating example for this construction is given by X being a subset of the unit sphere in $\mathbb{R}^3$. Then $\text{Con}(X)$ has its apex at the origin of $\mathbb{R}^3$ and consists of half-rays starting at 0_X which pass through X . Using spherical coordinates and the law of cosines, one can think of $d_c((x, t), (y, s))$ as being the length of the side opposite of 0_X in the triangle formed by these three points in $\mathbb{R}^3$.

4.2.5 Lemma (Elevator wins eventually). *Let A and B be two half-spaces in $\mathbb{R}^2$ and let $\alpha > 0$ be the angle measure of the dihedral angle $D = A \cap B$. Then the puff pastry $\mathcal{P}$ corresponding to the array*

$$\mathcal{A} = \underbrace{(A, B, A, B, \dots)}_{\lceil \frac{\pi}{\alpha} \rceil + 1 \text{ times}} \quad (4.2.3)$$

is end-to-end convex.

Proof. The trick to this proof is to view $\mathcal{P}$ as the cone over a certain gluing construction. First, observe that $\text{Con}(\mathbb{S}^1)$ is (isometric to) $\mathbb{R}^2$. $\mathcal{P}$ consists of several copies of $\mathbb{R}^2$ glued along A and B , alternating. Thus, $\mathcal{P}$ is isometric to the cone over the following construction X of $\lceil \frac{\pi}{\alpha} \rceil + 2$ glued unit circles.




The diagram should not be viewed as being embedded in the plane but rather as an abstract gluing of several unit circles. Each long arc on the left and on the right has length $\pi - \alpha$ and each short arc on the top and on the bottom has length α , so that each circle individually has length 2π . Each circle and its part of the cone correspond to a level in the puff pastry. Each arc on the left and on the right corresponds to A and B , alternating, with $A \cap B$ corresponding to the bottom arc shared by all circles. The arc on the top of length α is the part of each level that does belong only to that level, i.e. the complement of $A \cup B$. In each of these descriptions we of course have to include the part of the cone over the mentioned arcs.

Now the proof almost writes itself with this diagram. Observe that end-to-end convexity of $\mathcal{P}$ from this point of view means that any geodesic connecting points from the (part of the cone corresponding to the) innermost circle to the outermost circle stays entirely in the union of these two circles. When connecting points with different heights in the cone, the geodesic does all necessary “circle switching” continuously through its path. Loosely speaking, the geodesic can at most travel a length of π “in X -direction” since that is the diameter of the gluing construction and the conic height of the endpoints is irrelevant for deciding on which arc of length α the levels are traversed. Thus, we can further simplify end-to-end convexity of $\mathcal{P}$ to the following: any geodesic with length $< \pi$ connecting points from the innermost circle to the outermost circle on the same height (so both lie on the same conic section, which is just a diagram as above) stays entirely in the union of these two circles. This holds if the zigzag line, i.e. the concatenation of the arcs at the top excluding the innermost and outermost, has length at least π . But this line is formed by $\lceil \frac{\pi}{\alpha} \rceil$ arcs of length α , one for each circle minus the first and last, and so end-to-end convexity follows. $\square$

There is one more quick observation we have to make before the criterion, formulated in another lemma. We also recall from elementary geometry that a hyperplane A' is said to support a set $A \subseteq \mathbb{R}^n$, if A is contained in one of the half-spaces bounded by A' and $\partial(A) \cap A' \neq \emptyset$, i.e. A' contains a boundary point of A . This definition of course fits nicely together with 4.2.2(ii).


4.2.6 Lemma (Minimum is preserved in extending half-spaces). *Let A and B be two closed and convex subsets of $\mathbb{R}^n$ with $\text{Int}(A \cap B) \neq \emptyset$. Let $x, z \in \mathbb{R}^n$ and consider the function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, $f(y) := d(x, y) + d(y, z)$. Let $y_0 \in A \cap B$ be a minimum of f restricted to $A \cap B$. Then there exists half-spaces A' and B' containing A and B , respectively, such that x_0 is still a minimum of f restricted to $A' \cap B'$.*

Proof. Clearly, if $[x, z] \cap (A \cap B) \neq \emptyset$, then any point of this intersection is a global minimum of f and the choice of extending half-spaces is irrelevant. So we assume the segment $[x, z]$ does not intersect $A \cap B$. In this case, we conclude that y_0 must lie on the boundary of $A \cap B$. To see this it is very helpful to think of the level sets of f as ellipsoids with focal points p and q . But then clearly any ellipsoid that intersects $\text{Int}(A \cap B)$ can be shrunk to an ellipsoid that only touches $A \cap B$, which is again a closed and convex set, in a contact point. This also shows that y_0 is unique.



Note that $\partial(A \cap B) \subseteq \partial(A) \cup \partial(B)$ holds in any topological space, so we distinguish whether y_0 belongs to one or both boundaries. If $y_0 \in \partial(A) \setminus \partial(B)$, then we set A' to be the half-space corresponding to a supporting hyperplane H_1 at y_0 , such that y_0 is also a point of contact between H_1 and the ellipsoid E it lies on. Then y_0 is also a minimum of $f|_{A'}$ and we can consequently choose B' arbitrarily. The case $y_0 \in \partial(B) \setminus \partial(A)$ is symmetric. So finally, we consider $y_0 \in \partial(A) \cap \partial(B)$. Then we choose A' and B' such that they extend A and B , respectively, and the dihedral angle formed by them supports $A \cap B$ at y_0 . Then y_0 is still a minimum of $f|_{A' \cap B'}$, because it still is a contact point between E and $A' \cap B'$. $\square$

4.2.7 Remark (Nonempty interior is necessary). The following example shows that only requiring nonempty intersection of A and B is not sufficient to obtain extending half-spaces with the desired properties as in 4.2.6. Consider two disks with centers a and b , respectively, meeting in a contact point y . In particular, $y \in [a, b]$. Let $x = z$ be such that $[x, y]$ is perpendicular to $[a, b]$.



On the one hand, if one chooses half-spaces A' and B' corresponding to the lines tangent to A and B at x , respectively, then $A' \cap B'$ actually contains x and so y cannot be a minimum of $f|_{A' \cap B'}$. On the other hand, any other choice of extending half-spaces contains points closer than y to x in their intersection, since then either the intersection point of the corresponding hyperplanes is closer than y to x or $\text{Int}(A' \cap B')$ contains points closer than y to x . The success of the construction of the dihedral angle $A' \cap B'$ supporting $A \cap B$ at y_0 in 4.2.6 relies on the nonempty interior of $A \cap B$, which guaranties that the dihedral angle has nonzero angle measure.

The by now many times mentioned second criterion for end-to-end convexity is a generalization of 4.2.5 to higher dimensions and more general sets instead of half-spaces.

4.2.8 Proposition (Criterion for end-to-end convexity II). *Let A and B be closed and convex subsets of $\mathbb{R}^n$ such that any dihedral angle supporting $A \cap B$ at any point has angle measure at least $\alpha > 0$. Then the puff pastry $\mathcal{P}$ corresponding to the array*

$$\mathcal{A} = \underbrace{(A, B, A, B, \dots)}_{\lceil \frac{\pi}{\alpha} \rceil + 1 \text{ times}} \quad (4.2.4)$$

is end-to-end convex.

Proof. First, observe that the fact that any supporting dihedral angle has nonzero angle measure implies $\text{Int}(A \cap B) \neq \emptyset$, which justifies the use of 4.2.6 later on. Set $m := \lceil \frac{\pi}{\alpha} \rceil + 1$ and fix $x, z \in \mathbb{R}^n$. Consider the function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ defined by $f(p) := d(x, p) + d(p, z)$. Choose $y \in A \cap B$ such that y is a minimum of $f|_{A \cap B}$ and recall the notation introduced in 4.1.2. In order to prove that $\mathcal{P}$ is end-to-end convex, it suffices to show that the geodesic segment $[x^1, z^{m+1}]$ contains $y^1 = y^{m+1}$. Indeed, in this case the geodesic $[x^1, z^{m+1}]$ is the result of the concatenation of $[x^1, y^1]$ and $[y^{m+1}, z^{m+1}]$ and so stays entirely in $\mathcal{R}_1 \cup \mathcal{R}_{m+1}$.

As in 4.2.6, we can choose y such that it lies in $\partial(A \cap B)$. In particular, if $y \in \partial(A) \setminus \partial(B)$, then y is the contact point between A and the smallest ellipsoid E with focal points x and z that intersects A . Hence it follows that already in the puff pastry $P_1 := \mathbb{R}^n \sqcup_A \mathbb{R}^n$, the segment $[x^1, z^2]$ would meet $y^1 = y^2$. Moreover, the partial segment $[y^2, z^2] \subseteq [x^2, z^2]$ has a proper

subsegment entirely contained in B , viewed as a subset of the second level, since $y \in \text{Int}(B)$. This implies that adding an additional layer to P_1 glued along B does not change the geodesic connecting x in the first level and z in the last level. In other words, if $P_2 := P_1 \sqcup_B \mathbb{R}^n$, then $[x^1, z^2]_{\mathcal{P}_1} = [x^1, z^3]_{\mathcal{P}_2}$. This of course does not change when adding additional layers glued along A and B , and the case $y \in \partial(B) \setminus \partial(A)$ is analogous by switching the roles of A and B .

Thus, the only case left to check is $y \in \partial(A) \cap \partial(B)$. Let A' and B' be the half-spaces constructed in 4.2.6. Note that A' and B' admit elementary factorizations $A' = \mathbb{R}^{m-2} \times A''$, $B' = \mathbb{R}^{m-2} \times B''$, where A'' and B'' are half-spaces in $\mathbb{R}^2$. In an appropriately chosen basis, the half-space A' can be represented by the set of points whose last coordinate is larger than or equal to some $b \in \mathbb{R}$. Choose A'' to be the set of points whose second coordinate is larger than or equal to the same b (in an appropriate basis that can be extended to the basis in which A' has this representation). Similarly, the same process can be done with B' and B'' . Let $\mathcal{P}'$ be the puff pastry corresponding to the array

$$\mathcal{A}' = \underbrace{(A', B', A', B', \dots)}_{\lceil \frac{\pi}{\alpha} \rceil + 1 \text{ times}}. \quad (4.2.5)$$

Now $\mathcal{P}'$ admits a factorization $\mathcal{P}' = \mathbb{R}^{m-2} \times \mathcal{P}''$, where $\mathcal{P}''$ is the puff pastry corresponding to the array

$$\mathcal{A}'' = \underbrace{(A'', B'', A'', B'', \dots)}_{\lceil \frac{\pi}{\alpha} \rceil + 1 \text{ times}}. \quad (4.2.6)$$

By assumption, the dihedral angle $A' \cap B'$ has angle measure at least α and consequently the same holds for $A'' \cap B''$. Then 4.2.5 implies that $\mathcal{P}''$ is end-to-end convex. Moreover, since $\mathcal{P}' = \mathbb{R}^{m-2} \times \mathcal{P}''$, it follows that $\mathcal{P}'$ is end-to-end convex as well. In particular, the geodesic segment $[x^1, z^{m+1}]$ contains $y^1 = y^{m+1}$. Finally, we recall via 4.2.3 that the only possibility where $[x^1, z^{m+1}]$ does not stay in $\mathcal{R}_1 \cup \mathcal{R}_{m+1}$ is if it is shorter to jump back and forth between A and B in contrast to going directly via the intersection. Now already in $\mathcal{P}'$, it is shorter to go via the intersection, and the path that jumps back and forth between A and B is at least as long as the path jumping back and forth between A' and B' , cf. 4.1.5. Thus, we conclude that also $[x^1, z^{m+1}]$ goes through $y^1 = y^{m+1}$, and we are done. $\square$

There is one last proposition we need before we can prove the collision theorem. In a way it is yet another criterion for end-to-end convexity, but this time relying on the fact that the sets of the array have ε -wide corners.

4.2.9 Definition (ε -wide corners). *Let $\varepsilon > 0$.*

- (i) *Let A be a closed and convex subset of $\mathbb{R}^n$. We say that A has ε -wide corners if for all $p \in A$ an initial piece of a cone with apex p and aperture ε is contained in A . In other words, the maximum angle of any two generating lines of the cone is ε .*
- (ii) *Let $A_1, \dots, A_m$ be closed and convex subsets of $\mathbb{R}^n$. Assume that the intersection of each subset of $\{A_1, \dots, A_m\}$ has ε -wide corners, i.e. $\bigcap_{i \in S} A_i$ has ε -wide corners for all S in the power set of $\{1, \dots, m\}$. We abbreviate this fact by saying that all intersections of $A_1, \dots, A_m$ have ε -wide corners.*

4.2.10 Proposition (Wide corners array and end-to-end convexity). *Let $\varepsilon > 0$ and let $m \in \mathbb{N}$. Then there is a sequence $I_\varepsilon(m) = (i_1, \dots, i_M)$ of natural numbers depending on m and ε such that the following hold:*

- (i) $i_j \in \{1, \dots, m\}$ for all $j = 1, \dots, M$ and for all $k = 1, \dots, m$ there is a j such that $i_j = k$, i.e. every number from 1 to m appears at least once in this sequence.
- (ii) If $A_1, \dots, A_m$ is any collection of closed and convex subsets of $\mathbb{R}^n$ with $\bigcap_{i=1}^m A_i \neq \emptyset$ whose intersections all have ε -wide corners, then the puff pastry $\mathcal{P}$ corresponding to the array $\mathcal{A} = (A_{i_1}, \dots, A_{i_M})$ is end-to-end convex.

Proof. Before we begin with the proof, we want to emphasize that the choice of the subsets does not matter for the sequence, i.e. once we have constructed $I_\varepsilon(m)$, any puff pastry for any array indexed by this sequence is end-to-end convex! Also, there is some interest in finding the smallest M such that the statement holds, since (puff pastries corresponding to) larger arrays created from (the arrays of) already end-to-end convex puff pastry in the sense of 4.2.1 are also end-to-end convex. Indeed, we will see that we can always assume that $M \leq (\lceil \frac{\pi}{\varepsilon} \rceil + 1)^m$.

We construct $I_\varepsilon(m)$ inductively with respect to m . If $m = 1$, there is nothing to do since any puff pastry with two layers is end-to-end convex by definition, so $I_\varepsilon(1) = (1)$ satisfies the assumptions. Suppose we have already constructed $I_\varepsilon(m)$. We claim that $I_\varepsilon(m+1)$ can be constructed from $I_\varepsilon(m)$ by replacing each occurrence of m in $I_\varepsilon(m)$ by the alternating sequence

$$\underbrace{m, m+1, m, m+1, \dots}_{\lceil \frac{\pi}{\varepsilon} \rceil + 1 \text{ times}} \quad (4.2.7)$$

Clearly, $I_\varepsilon(m+1)$ still satisfies (i), for $m+1$ instead of m , of course. Let $A_1, \dots, A_{m+1}$ be any collection of closed and convex sets satisfying the conditions in (ii). Let $\mathcal{A}$ be an array for $A_1, \dots, A_{m+1}$ indexed by $I_\varepsilon(m+1)$ and let $\mathcal{P}$ be the corresponding puff pastry. We have to show that $\mathcal{P}$ is end-to-end convex. To this end let $\mathcal{P}'$ be the puff pastry for an array $\mathcal{A}'$ which is constructed from $\mathcal{A}$ in the following way: each alternating sequence

$$\underbrace{A_m, A_{m+1}, A_m, A_{m+1}, \dots}_{\lceil \frac{\pi}{\varepsilon} \rceil + 1 \text{ times}} \quad (4.2.8)$$

in $\mathcal{A}$ that resulted from the construction of $I_\varepsilon(m+1)$ out of $I_\varepsilon(m)$ is replaced by the sequence

$$\underbrace{A_m \cap A_{m+1}, A_m \cap A_{m+1}, \dots}_{\lceil \frac{\pi}{\varepsilon} \rceil + 1 \text{ times}} \quad (4.2.9)$$

If we set $B := A_m \cap A_{m+1}$, then we see that $\mathcal{A}'$ is an array where only the m closed and convex sets $A_1, \dots, A_{m-1}, B$ appear. Moreover, $A_1, \dots, A_{m-1}, B$ satisfy the conditions in (ii), and so the puff pastry $\mathcal{P}''$ corresponding to the array $\mathcal{A}''$ for $A_1, \dots, A_{m-1}, B$ indexed by $I_\varepsilon(m)$ is end-to-end convex. From the way that $I_\varepsilon(m+1)$ is constructed out of $I_\varepsilon(m)$, we observe that $\mathcal{A}'$ and $\mathcal{A}''$ satisfy the requirements of 4.2.1, so we can apply the criterion to $\mathcal{P}''$ and $\mathcal{P}'$ and conclude that $\mathcal{P}'$ is end-to-end convex as well. Now the only potential problem left for the end-to-end convexity of $\mathcal{P}$ is if in the alternating sequence $A_m, A_{m+1}, \dots$ it is shorter for a geodesic to jump around through each level between A_m and A_{m+1} compared to going via $B = A_m \cap A_{m+1}$. Since any intersection of $A_1, \dots, A_{m+1}$ has ε -wide corners, this holds in particular for $A_m \cap A_{m+1}$, which implies that any dihedral angle supporting $A_m \cap A_{m+1}$ at any of its points has angle measure at least ε . So by 4.2.8, the puff pastry $\mathcal{P}'''$ corresponding to the array

$$\mathcal{A}''' := \left(\underbrace{A_m, A_{m+1}, A_m, A_{m+1}, \dots}_{\lceil \frac{\pi}{\varepsilon} \rceil + 1 \text{ times}} \right) \quad (4.2.10)$$

is end-to-end convex. But this solves the only problem where something could have gone wrong, since $\mathcal{A}$ and $\mathcal{A}'$ agree everywhere else but these sequences! In other words, since $\mathcal{P}'$ is end-to-end convex and $\mathcal{A}$ and $\mathcal{A}'$ differ only in the sequences we constructed and this sequence in $\mathcal{A}$ viewed as an individual array $\mathcal{A}'''$ results in an end-to-end convex puff pastry $\mathcal{P}'''$, it follows that $\mathcal{P}$ is end-to-end convex as well. Finally, we note that the upper bound on M follows directly from the construction. $\square$


The above proposition can be regarded as a generalization of 4.2.8, since setting $m = 2$ in 4.2.10 gives precisely the construction from 4.2.8.

4.3 The collision theorem

Now we are in a position to prove the collision theorem. At first, we define the setting of the theorem.

4.3.1 Definition (Building a billiard table). *Let $A_1, \dots, A_m$ be closed and convex subsets of $\mathbb{R}^n$ such that $\partial(A_i)$ is a smooth hypersurface. So for example, if $n = 2$, then neither a (filled) rectangle nor a straight line would be valid choices for A_i .*

- (i) *The closure of the complement of the union of such sets is called a billiard table T , i.e. $T := \text{Cl}(\mathbb{R}^n \setminus (\bigcup_{i=1}^m A_i))$.*
- (ii) *Each set A_i is called a wall of the table T .*
- (iii) *A billiard trajectory is a unit-speed piecewise linear path that only changes its direction when meeting $\partial(A_i)$ and does so always following the standard laws of physics at the break points, i.e. the angle of reflection is equal to the angle of incidence. A trajectory can, but need not be defined on all of $\mathbb{R}$.*
- (iv) *The break points of the trajectory are called collisions.*



For simplicity, we assume that each trajectory only meets one wall at a time. The assumption on $\partial(A_i)$ being a smooth hypersurface guaranties on the one hand that each A_i has the correct “dimension”, and on the other hand the smoothness ensures that there are no problems with the break points of trajectories, e.g. it would be impossible to calculate the angle if the trajectory hits the corner of a rectangle.

4.3.2 Theorem (Collision theorem). *Let T be a billiard table in $\mathbb{R}^n$ with m walls such that the walls satisfy $\text{Int}(\bigcap_{i=1}^m A_i) \neq \emptyset$ and all their intersections have ε -wide corners. Then the number of collisions of any billiard trajectory is bounded from above by $M := (\lceil \frac{\pi}{\varepsilon} \rceil + 1)^{m^2}$.*

Proof. We prove the theorem by induction on m . If $m = 1$, then any trajectory that meets $\partial(A_1)$ can never return, since the angle of incidence has to agree with the angle of reflection and A_1 is convex. So every trajectory has at most one collision.

Suppose we have shown the statement for any billiard table with $m - 1$ walls. Let T be a table with m walls and let $\gamma : \mathbb{R} \rightarrow T$ be a trajectory that meets the walls in the order $A_{i_1}, \dots, A_{i_N}$. Denote by a_j the time of the j -th collision point of γ , so $\gamma(a_j) \in \partial(A_{i_j})$ for all $j = 1, \dots, N$. Let $\mathcal{A}$ be the array that contains all sets with which γ collides in order, i.e. $\mathcal{A} := (A_{i_1}, \dots, A_{i_N})$. By induction hypothesis, any number of $M' := (\lceil \frac{\pi}{\varepsilon} \rceil + 1)^{(m-1)^2} + 1$ consecutive elements in $\mathcal{A}$ has to contain each element of $\{A_1, \dots, A_m\}$ at least once. Indeed, otherwise (a suitable restriction of) γ would violate the bound on the number of collisions in some table T' with less than m walls. Let $\mathcal{P}$ be the puff pastry for the array $\mathcal{A}$. Let $\bar{\gamma} : \mathbb{R} \rightarrow \mathcal{P}$ be the curve defined by

$$\bar{\gamma}(t) := \begin{cases} \gamma(t)^1, & \text{if } t < a_1 \\ \gamma(t)^j, & \text{if } t \in [a_{j-1}, a_j] \\ \gamma(t)^{N+1}, & \text{if } t > a_N. \end{cases}$$

In other words, $\bar{\gamma}$ can be viewed as a lifting of γ to $\mathcal{P}$ in the sense that it starts out on the first level of $\mathcal{P}$ until it hits A_1 , then jumps to the second level and stays there until it hits A_2 , and so on. Since the angle of incidence is equal to the angle of reflection at each break point, $\bar{\gamma}$ is a local geodesic in $\mathcal{P}$. In other words, the shortest path from $x^j \in [\gamma(a_{j-1})^j, \gamma(a_j)^j]$ to $y^{j+1} \in [\gamma(a_j)^{j+1}, \gamma(a_{j+1})^{j+1}]$ always goes through $\gamma(a_j)^j = \gamma(a_j)^{j+1}$, i.e. this shortest path is just the concatenation $[x^j, \gamma(a_j)^j] \cdot [\gamma(a_j)^{j+1}, y^{j+1}]$. This is the case because $\gamma(a_j)$ is precisely the contact point in $\mathbb{R}^n$ of A_j and the ellipsoid with focal points x and y , cf. 4.2.8, which is guaranteed by the equality of the angle of incidence and the angle of reflection. Now $\mathcal{P}$ is a CAT(0) space, so every local geodesic is a geodesic by 3.1.5. Since γ does not meet $\bigcap_{i=1}^m A_i$, neither does the geodesic $\bar{\gamma}$ and so its image does not lie in $R_1 \cup R_{m+1}$, which in particular implies that $\mathcal{P}$ is not end-to-end convex. Let $I_\varepsilon(m) = (k_1, \dots, k_{N'})$ be the sequence constructed in 4.2.10. Then the array $\mathcal{A}' := (A_{k_1}, \dots, A_{k_{N'}})$ results in an end-to-end convex puff pastry $\mathcal{P}'$.

We now want to show that if the number of collisions N of γ is large enough, $\mathcal{A}$ can be obtained from $\mathcal{A}'$ by inserting additional copies of the A_i 's in the sense of 4.2.1, which would imply that $\mathcal{P}$ is also end-to-end convex, a contradiction. Since they already consist of the same sets, we only have to verify that the initial order of the elements of $\mathcal{A}'$ is preserved in $\mathcal{A}$, which follows from a purely combinatorial argument. Indeed, this is always the case if there are at least N' distinct sequences of length M' in $\mathcal{A}$: recalling that by induction hypothesis any sequence of length M' contains each set in $\{A_1, \dots, A_m\}$ at least once, we can view the occurrence of A_{k_1} in the first sequence, the occurrence of A_{k_2} in the second sequence and so on up to the occurrence of $A_{k_{N'}}$ in the N' -th sequence as the original structure of $\mathcal{A}'$ onto which we add additional copies of the

sets in $\{A_1, \dots, A_m\}$ to obtain $\mathcal{A}$. Thus, if $N \geq M' \cdot N'$, this contradictorily implies that $\mathcal{P}$ is end-to-end convex. In summary, any trajectory in T can have at most $M' \cdot N'$ collisions. Using 4.2.10 and the induction hypothesis, we compute

$$M' \cdot N' \leq ((\lceil \frac{\pi}{\varepsilon} \rceil + 1)^{(m-1)^2} + 1)(\lceil \frac{\pi}{\varepsilon} \rceil + 1)^m \leq (\lceil \frac{\pi}{\varepsilon} \rceil + 1)^{m^2} = M, \quad (4.3.1)$$

and the claim follows. $\square$

Intuitively, the existence of such a bound may not be a big surprise. However, it certainly is remarkable that the maximum number of collisions is independent of the dimension of the billiard table as well as the specific shape of the walls! Only the number of walls and the width of their corners is relevant.

To conclude this chapter, we give an interesting corollary of the collision theorem to the so called system of hard balls. This is the setting hinted at in the description of this chapter, somewhat resembling a real game of billiards. Consider m homogeneous balls $B_1, \dots, B_m$ in $\mathbb{R}^3$ with radii R_i and evenly distributed masses M_i . Suppose they travel along straight lines with constant speed. We suppose that at each point in time there occurs at most one collision and each collision happens between two balls only. If two balls collide, their new directions and velocities are calculated using the laws of classical mechanics. Such a constellation is called a system of m hard balls. If considering the same situation in $\mathbb{R}^2$ and assuming equal radii and masses, then this corresponds to a strike on a “real” (albeit infinite) billiard table. For this corollary we need a lemma that ensures the existence of wide corners in convex sets with some additional properties. To this end we introduce some notions and theorems from convex geometry, which we will not prove.

4.3.3 Lemma (Compactness and wide corners). *Let A be a compact and convex subset of $\mathbb{R}^n$ with $\text{Int}(A) \neq \emptyset$. Then there exists $\varepsilon > 0$ such that A has ε -wide corners.*

Proof. The condition is trivially satisfied for any interior point, so it is only necessary to check the condition for boundary points. Let $p \in \partial(A)$, $q \in \text{Int}(A)$ and let $r > 0$ be such that $B_r(q) \subseteq A$. Since A is closed and convex, the segment $[x, p]$ is contained in A for each $x \in B_r(q)$. Now these segments together already form (part of) an m -dimensional cone with apex p , since $B_r(q)$ is m -dimensional. So every point p has a cone with an angle of aperture δ for some $\delta > 0$ contained in A . Finally, we have to show that these apertures do not become arbitrarily small. To this end, fix $q \in \text{Int}(A)$ and $r > 0$ such that $B_r(q) \subseteq A$ and define the function $f : \partial(A) \rightarrow \mathbb{R}$ by $f(p) := \sup_{x, y \in B_r(q)} \angle_p(x, y)$, i.e. the largest angle between any two segments emanating from p into $B_r(q)$. This is precisely the angle of the aperture of the cone with apex p and base $B_r(q)$. In particular, f is continuous since the angle is continuous and $B_r(q)$ is bounded. Additionally, we have $f(p) > 0$ for all $p \in \partial(A)$ by the arguments above. Since A is compact, $\partial(A)$ is compact as well, so f attains its infimum and we can define $\varepsilon := \min_{p \in \partial(A)} f(p) > 0$. $\square$

The construction from convex geometry we need is called the recession cone. Intuitively, in an unbounded convex set this describes the directions in which the set “continues forever”.

4.3.4 Definition (Recession cone). *Let $A \neq \emptyset$ be a convex subset of $\mathbb{R}^n$ and let $y \in \mathbb{R}^n$.*

- (i) *A is said to recede in direction y , if $x + ty \in A$ for all $x \in A$ and $t > 0$. In other words, the half-ray emanating at x that extends the segment $[x, y]$ is contained in A for all x in A .*
- (ii) *The recession cone of A is the set of all directions y in which A recedes and is denoted by $\text{RC}(A)$. We include $0 \in \mathbb{R}^n$ as a trivial direction in which any set recedes.*

Note that if a convex set A recedes in direction y , it does so in direction sy for all $s > 0$ as well. In other words, any positive multiple of a receding direction is again a receding direction. Moreover, if x and y are two receding directions, then so is $x + y$, so $\text{RC}(A)$ has some similarities to a linear subspace. Some basic properties of the recession cone are collected in the following proposition, cf. [Roc70, Thms. 8.1, 8.2 & 8.4].

4.3.5 Proposition (Properties of the recession cone). *Let A be a convex subset of $\mathbb{R}^n$. Then we have:*

- (i) $\text{RC}(A)$ is a convex subset of $\mathbb{R}^n$.
- (ii) If A is closed, so is $\text{RC}(A)$.
- (iii) A is bounded if and only if $\text{RC}(A) = \{0\}$. □

The sets for which we want to show the existence of ε -wide corners have a common center of symmetry, where we understand a point p of a set A to be a center of symmetry if A is symmetric with respect to reflections through p . Before we continue, we recall another definition from elementary geometry.

4.3.6 Definition (Relative interior). *Let A be a subset of $\mathbb{R}^n$.*

- (i) *The affine hull of A , denoted by $\text{Aff}(A)$, is the set of all affine combinations of elements of A , i.e.*

$$\text{Aff}(A) := \left\{ \sum_{i=1}^k t_i a_i \mid k \in \mathbb{N}, a_i \in A, t_i \in \mathbb{R}, \sum_{i=1}^k t_i = 1 \right\}. \quad (4.3.2)$$

- (ii) *The relative interior of A is denoted by $\text{RInt}(A)$ and is defined as the interior of A viewed as a set in the subspace $\text{Aff}(A) \subseteq \mathbb{R}^n$, that is,*

$$\text{RInt}(A) := \text{Int}_{\text{Aff}(A)}(A) = \{a \in A \mid \exists \varepsilon > 0 : B_\varepsilon(a) \cap \text{Aff}(A) \subseteq A\}. \quad (4.3.3)$$

The relative interior is useful when studying lower dimensional objects in higher dimensional spaces. As an easy example, the interior of a (closed) square in $\mathbb{R}^3$ is obviously empty, but its relative interior is the usual interior one obtains when viewing it as a subset of $\mathbb{R}^2$. The following proposition corresponds to [Roc70, Thm. 8.3].

4.3.7 Proposition (Any ray belongs to the cone). *Let A be a closed and convex subset of $\mathbb{R}^n$ and let $0 \neq y \in \mathbb{R}^n$. If there is an $x \in A$ such that $\{x + ty \mid t \geq 0\} \subseteq A$, then $y \in \text{RC}(A)$, i.e. for all $x \in A$ the set $\{x + ty \mid t \geq 0\}$ is contained in A . In particular, if $x \in \text{RInt}(A)$, then $\{x + ty \mid t \geq 0\} \subseteq \text{RInt}(A)$, so $y \in \text{RC}(\text{RInt}(A))$. □*

4.3.8 Lemma (Wide corners of symmetric sets). *Let $A_1, \dots, A_m$ be closed and convex subsets in $\mathbb{R}^n$ with nonempty interior and a common center of symmetry. Then there exists $\varepsilon > 0$ such that all their intersections have ε -wide corners.*

Proof. The main idea of this proof is to show that any closed and convex set with nonempty interior and a center of symmetry already has ε -wide corners for some $\varepsilon > 0$. Then this of course holds for all possible intersections by assumption, since all defining properties are preserved under finite intersections. As there are only finitely many intersections to consider, we choose the

minimal ε , since if a set has ε -wide corners, it clearly also has ε' -wide corners for all $\varepsilon' \leq \varepsilon$.

Let $A \subseteq \mathbb{R}^n$ be closed and convex, $\text{Int}(A) \neq \emptyset$ and without loss of generality let the origin be a center of symmetry of A , since otherwise we consider a suitable translation of A . Consider the recession cone $\text{RC}(A)$ of A . By 4.3.7 and A being symmetric around 0, it holds that if $y \in \text{RC}(A)$, then $-y \in \text{RC}(A)$. In particular, $\text{RC}(A)$ is a linear subspace and $x + \text{RC}(A) := \{x + p \mid p \in \text{RC}(A)\}$ is contained in A for all $x \in A$. We now distinguish the number of linearly independent vectors in $\text{RC}(A)$: If there are no linearly independent vectors in $\text{RC}(A)$, then $\text{RC}(A) = \{0\}$, so A is bounded by 4.3.5(iii). Since A is closed by assumption, A is compact and the claim follows by 4.3.3.

If $\text{RC}(A)$ contains one linearly independent vector, then $\text{RC}(A)$ is a line. By 4.3.7, we have $x + \text{RC}(A) \subseteq A$ for all $x \in A$, so A is a union of lines. Since A is convex, any segment between points belonging to distinct lines is contained in A . In other words, if A' is a hyperplane which contains 0 and is perpendicular to $\text{RC}(A)$, then the intersection $B := A \cap A'$ is again convex. In particular, B determines A completely, since B contains the foot point of every line in A . Moreover, B is compact since otherwise $\text{RC}(A)$ would contain another linearly independent vector by 4.3.7 and $\text{RInt}(B) \neq \emptyset$ since $\text{Int}(A) \neq \emptyset$. So we can apply 4.3.3 to B (viewed as a subset in $\mathbb{R}^2$) to obtain 2-dimensional cones with positive aperture bounded below for every point in B . Now these cones can be translated along $\text{RC}(A)$ to obtain the desired n -dimensional cones for every point in A .

This argument can readily be generalized to higher dimensions. Suppose there are $k < n$ linearly independent vectors in $\text{RC}(A)$, then $\text{RC}(A)$ is a k -dimensional linear subspace. By 4.3.7, we have $x + \text{RC}(A) \subseteq A$ for all $x \in A$, so A is a union of k -dimensional linear subspaces. Analogously to the case above, let A' be an $n - k$ -dimensional subspace normal to $\text{RC}(A)$. Then the intersection $B := A \cap A'$ uniquely determines A , and B is compact since otherwise $\text{RC}(A)$ would contain another linearly independent vector. Moreover, $\text{RInt}(B) \neq \emptyset$ and so we can apply 4.3.3 to B (viewed as a subset in $\mathbb{R}^k$) to obtain k -dimensional cones with positive aperture bounded below for every point in B . Again, these cones can be translated along $\text{RC}(A)$ to obtain the desired n -dimensional cones for every point in A . Finally, we note that if $\text{RC}(A)$ contains n linearly independent vectors, then $\text{RC}(A) = A = \mathbb{R}^n$. Thus, the statement holds in every case and we are done. $\square$

4.3.9 Corollary (Bound on collisions in hard ball system). *There exists an upper bound on the number of collisions in any hard ball system depending only on the number of balls, their radii and their masses. That is, this bound is independent of the initial directions and velocities of the balls.*

Proof. Let $M_i, R_i \in \mathbb{R}, a_i \in \mathbb{R}^3$ be the mass, radius and center, respectively, of the ball B_i , $i = 1, \dots, m$. Collect all the centers in a point p in $\mathbb{R}^{3m}$, so $p = (a_1, \dots, a_m)$. Clearly, if $d_{\mathbb{R}^3}(a_i, a_j) \leq R_i + R_j$, then $B_i \cap B_j \neq \emptyset$. The system of these inequalities for all i and j defines $\frac{n \cdot (n-1)}{2}$ “cylinders” $C_{i,j} := \{(x_1, \dots, x_m) \in \mathbb{R}^{3m} \mid d_{\mathbb{R}^3}(x_i, x_j) \leq R_i + R_j\}$. Each $C_{i,j}$ is closed and convex and contains a neighbourhood of $0 \in \mathbb{R}^{3m}$, so $\text{Int}(\bigcap_{i < j} C_{i,j}) \neq \emptyset$. Moreover, $\partial(C_{i,j})$ is a smooth hypersurface and so $T := \text{Cl}(\mathbb{R}^{3m} \setminus \bigcup_{i < j} C_{i,j})$ is a billiard table, with each point in T corresponding to a valid position of the hard ball system. Depending on the initial directions and velocities of the balls, the evolution of the hard ball system is described by the motion of the point $p \in \mathbb{R}^{3m}$, moving along a straight line until it meets one of the cylinders. Define a norm on $\mathbb{R}^{3m}$ via $\|q\|_T := \sqrt{\sum_{i=1}^m M_i \|q_i\|_{\mathbb{R}^3}^2}$, where $q = (q_1, \dots, q_m)$ and $q_i \in \mathbb{R}^3$. In the metric induced by this norm, the collisions of p with the sets $C_{i,j}$ follow the standard laws of physics,

i.e. the angle of reflection is equal to the angle of incidence. Thus, the estimate on the number of collisions in the hard ball system corresponds to the number of collisions of an appropriate trajectory in T . By 4.3.8, all intersections of the $C_{i,j}$'s have ε -wide corners for some ε depending on the radii and masses of the balls. This allows us to apply 4.3.2. $\square$

Note that this corollary can easily be generalized to arbitrary dimensions.

References

- [AKP19a] S. B. Alexander, V. Kapovitch, and A. Petrunin. Alexandrov geometry. arXiv:1903.08539v1, 2019.
- [AKP19b] S. B. Alexander, V. Kapovitch, and A. Petrunin. *An Invitation to Alexandrov Geometry: CAT(0) spaces*. SpringerBriefs in Mathematics. Springer, Berlin, 2019.
- [Bal95] W. Ballman. *Lectures on spaces of nonpositive curvature*, volume 25 of *DMV-Seminars*. Birkhäuser, Basel, 1995.
- [BBI01] D. Burago, Y. Burago, and S. Ivanov. *A course in metric geometry*, volume 33 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2001.
- [BFK98a] D. Burago, S. Ferleger, and A. Kononenko. A geometric approach to semi-dispersing billiards. *Ergodic Theory and Dynamical Systems*, 18(2):303–319, 1998.
- [BFK98b] D. Burago, S. Ferleger, and A. Kononenko. Uniform estimates on the number of collisions in semi-dispersing billiards. *Annals of Mathematics, Second Series*, 147(3):695–708, 1998.
- [BH99] M. R. Bridson and A. Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Comprehensive Studies in Mathematics*. Springer, Berlin, 1999.
- [Hol13] I. Holopainen. Metric geometry. <https://wiki.helsinki.fi/display/mathstatKurssit/Metric+geometry,+fall+2013?preview=/112441203/119537966/MetGeo.pdf>, 2013.
- [Jos97] J. Jost. *Nonpositive curvature: geometric and analytic aspects*. Lectures in Mathematics ETH Zürich. Birkhäuser, Basel, 1997.
- [Lan04] U. Lang. Length spaces. <https://people.math.ethz.ch/~lang/LengthSpaces.pdf>, 2004.
- [Pap13] A. Papadopoulos. *Metric spaces, convexity and nonpositive curvature*, volume 6 of *IRMA Lectures in Mathematics and Theoretical Physics*. European Mathematical Society, 2nd edition, 2013.
- [Pau04] G. Paulik. *Gluing spaces and Analysis*. PhD thesis, University of Bonn, 2004.
- [Pet19] A. Petrunin. <https://mathoverflow.net/questions/346624>, 2019.
- [Roc70] R. T. Rockafellar. *Convex analysis*, volume 28 of *Princeton Mathematical Series*. Princeton University Press, 1970.

- [Shi93] K. Shiohama. *An introduction to the geometry of Alexandrov spaces*. Research Institute of Mathematics, Global Analysis Research Center, Seoul National University, 1993.
- [Stu06] K. T. Sturm. On the geometry of metric measure spaces. *Acta Mathematica*, 196(1):65–131, 2006.
- [Val64] F. Valentine. *Convex sets*. McGraw-Hill, New York City, NY, 1964.