Chapter 1

Model plane

A. Trigonometry

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

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

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

Let $\varpi^\kappa = \operatorname{diam} \mathbb{M}^2(\kappa)$, so $\varpi^\kappa = \infty$ if $\kappa \leqslant 0$ and $\varpi^\kappa = \pi/\sqrt{\kappa}$ if $\kappa > 0$; ϖ is just a cursive form of π.

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

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

The angle of $[xyz]$ at x will be denoted by $\measuredangle\left[x\,{}^{y}_{z}\right]$.

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

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

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

Trigonometric functions. We will need three *trigonometric functions* in $\mathbb{M}^2(\kappa)$: *cosine*, *sine*, and *modified distance*, denoted by cs^κ, sn^κ, and md^κ, respectively.

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

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

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

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

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

$$\operatorname{sn}^{\pm\kappa}=\tfrac{1}{\sqrt{\kappa}}\cdot\operatorname{sn}^{\pm1}(x\cdot\sqrt{\kappa});\quad \operatorname{cs}^{\pm\kappa}=\operatorname{cs}^{\pm1}(x\cdot\sqrt{\kappa});\quad \operatorname{md}^{\pm\kappa}=\tfrac{1}{\kappa}\cdot\operatorname{md}^{\pm1}(x\cdot\sqrt{\kappa});$$

$$\operatorname{sn}^{-1}x=\sinh x;\quad \operatorname{cs}^{-1}x=\cosh x;\quad \operatorname{md}^{-1}x=\cosh x-1;$$

$$\operatorname{sn}^{0}x=x;\quad \operatorname{cs}^{0}x=1;\quad \operatorname{md}^{0}x=\tfrac12\cdot x^2;$$

$$\operatorname{sn}^{1}x=\sin x;\quad \operatorname{cs}^{1}x=\cos x;\quad \operatorname{md}^{1}x=\begin{cases}1-\cos x & \text{for } x\leqslant\pi,\\ 2 & \text{for } x>\pi.\end{cases}$$

Note that

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

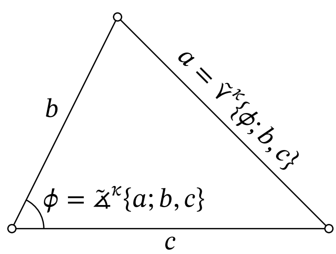

Let ϕ be the angle of $\tilde{\triangle}^\kappa\{a,b,c\}$ opposite to a. In this case, we will write

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

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

$$\tilde{\curlyvee}^\kappa\{\phi;b,-c\}=\tilde{\curlyvee}^\kappa\{\phi;-b,c\}\coloneqq\tilde{\curlyvee}^\kappa\{\pi-\phi;b,c\};$$

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

1.1. Properties of standard functions.

(a) *For fixed a and ϕ, the function $y(t) = \operatorname{md}^\kappa\left(\tilde{\gamma}^\kappa\{\phi; a, t\}\right)$ satisfies the differential equation*

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

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

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

for $y \neq 0$.

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

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

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

$$\phi \mapsto \tilde{\gamma}^\kappa\{\phi; b, c\}$$

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

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

$$\kappa \mapsto \tilde{\measuredangle}^\kappa\{a; b, c\} \quad \text{and} \quad \kappa \mapsto \tilde{\gamma}^\kappa\{\phi; b, c\}$$

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

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

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

Then they have the same sign.

All the properties except (e), Alexandrov's lemma, can be shown by direct calculation. Alexandrov's lemma is reformulated in 6.3 and is proved there.

Cosine law. The formulas $a = \tilde{\gamma}^\kappa\{\phi; b, c\}$ and $\phi = \tilde{\measuredangle}^\kappa\{a; b, c\}$ can be rewritten using the cosine law in $\mathbb{M}^2(\kappa)$:

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

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

B. Hemisphere lemma

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

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

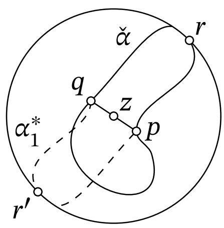

Assume $\ell < \pi$. Let $\check\alpha$ be a subarc of α of length ℓ, with endpoints p and q. Since $|p-q| \leqslant \ell < \pi$, there is a unique geodesic $[pq]$ in $\mathbb{S}^2$. Let z be the midpoint of $[pq]$. We claim that α lies in the open hemisphere centered at z. If not, α intersects the boundary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

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

a contradiction.

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

1.3. Exercise. Give a proof of 1.2 based on Crofton's formula.

Chapter 2

Metric spaces

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

A. Metrics and their relatives

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

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

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

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

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

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

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

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

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

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

will be called the *distance function* from x.

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

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

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

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

A map between two metric spaces is called an *isometry* if it is a bijection and preserves distances between points.

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

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

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

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

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

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

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

Now let us give examples showing that vanishing and infinite distance between distinct points can appear naturally and useful in constructions.

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

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

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

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

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

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

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

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

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

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

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

B. Notations

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

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

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

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

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

respectively.

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

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

respectively.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

(1) For some (and therefore any) point $p \in \mathcal{X}$ and any $R < \infty$, the closed ball $\overline{\mathrm{B}}[p, R] \subset \mathcal{X}$ is compact.

(2) The function $\operatorname{dist}_p : \mathcal{X} \to \mathbb{R}$ is proper for some (and therefore any) point $p \in \mathcal{X}$.

We will also often use the following two classical statements:

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

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

(a) *$\mathcal{X}$ is compact;*

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

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

(d) *$\mathcal{X}$ is complete and for any $\varepsilon > 0$ there is a* compact ε-net *in $\mathcal{X}$; that is,* $\mathrm{B}(K, \varepsilon) = \mathcal{X}$ *for a compact set $K \subset \mathcal{X}$.*

C. Length spaces

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

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

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

The following lemma is an easy exercise.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

2.8. Example. Consider a metric space $\mathcal{X}$ obtained by gluing a countable collection of disjoint intervals $\mathbb{I}_n$ of length $1+\frac{1}{n}$ where for each $\mathbb{I}_n$ one end is glued to p and the other to q. Then $\mathcal{X}$ carries a natural complete length metric such that $|p-q| = 1$, but there is no geodesic connecting p to q.

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

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

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

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

2.11. Exercise. Let $\mathcal{X}$ be a complete length space. Show that for any compact subset K in $\mathcal{X}$ there is a compact path-connected subset K' that contains K.

2.12. Definition. Consider two points x and y in a metric space $\mathcal{X}$.

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

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

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

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

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

The following lemma was essentially proved by Karl Menger [**116**, Section 6].

2.13. Lemma. *Let $\mathcal{X}$ be a complete metric space.*

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

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

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

Proof. Fix a pair of points $x, y \in \mathcal{X}$. Let $\varepsilon_n = \frac{\varepsilon}{2^{2\cdot n}}$.

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

This way we define $\alpha(t)$ for all dyadic rationals t in $[0,1]$. If $t \in [0,1]$ is not a dyadic rational, consider a sequence of dyadic rationals $t_n \to t$ as $n \to \infty$. By completeness of $\mathcal{X}$, the sequence $\alpha(t_n)$ converges; let $\alpha(t)$ be its limit. It is easy to see that $\alpha(t)$ does not depend on the choice of the sequence t_n, and $\alpha\colon [0,1] \to \mathcal{X}$ is a path from x to y. Moreover,

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

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

To prove (b), one should repeat the same argument taking midpoints instead of ε_n-midpoints. In this case, (1) holds for $\varepsilon_n = \varepsilon = 0$.

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

2.14. Exercise. Suppose that for any two distinct points x and y in a complete metric space $\mathcal{X}$ there is yet another point z such that

$$|x - z| + |y - z| = |x - y|.$$

Show that $\mathcal{X}$ is geodesic.

Since in a compact set a sequence of $\frac{1}{n}$-midpoints (z_n) contains a convergent subsequence, Lemma 2.13 implies the following.

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

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

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

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

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

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

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

Indeed, $\mathcal{X}$ is a length space; therefore for any $\varepsilon > 0$, the set $\overline{\mathrm{B}}[x, \rho(x) - \varepsilon]$ is a compact ε-net in B. Since B is closed and hence complete, it is compact by Proposition 2.4. △

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

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

Set $\varepsilon = \min_{y \in B}\{\rho(y)\}$; the minimum is defined since B is compact. From (2), we have $\varepsilon > 0$.

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

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

D. Convex sets

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

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

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

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

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

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

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

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

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

2.19. Proposition. *Let $\mathcal{X}_n$, be a sequence of geodesic spaces. Let ω be an ultrafilter on $\mathbb{N}$ (see Definition 4.1). Assume that $A_n \subset \mathcal{X}_n$ is a sequence of weakly convex sets, $\mathcal{X}_n \to \mathcal{X}_\omega$, and $A_n \to A_\omega \subset \mathcal{X}_\omega$ as $n \to \omega$. Then A_ω is a weakly convex set of $\mathcal{X}_\omega$.*

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

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

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

It remains to observe that α_ω is a geodesic path that lies in A_ω. ☐

E. Quotient spaces

Quotient spaces. Assume $\mathcal{X}$ is a metric space with an equivalence relation $\sim$. Note that given a family of pseudometrics ρ_α on $\mathcal{X}/\sim$, their least upper bound

$$\rho(x,y) = \sup_\alpha \{\rho_\alpha(x,y)\}$$

is also a pseudometric. If the projections $\mathcal{X} \to (\mathcal{X}/\sim, \rho_\alpha)$ are *short* (that is, *distance nonincreasing*), then so is $\mathcal{X} \to (\mathcal{X}/\sim, \rho)$.

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

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

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

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

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

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

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

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

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

is also short.

Note that

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

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

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

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

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

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

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

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

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

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

2.21. Proposition. *Assume $\mathcal{X}$ is a metric space and a group G acts on $\mathcal{X}$ by isometries. Then the projection $\pi\colon\ \mathcal{X} \to \mathcal{X}/G$ is a* submetry*; that is,* $\pi(\mathrm{B}(p,r)) = \mathrm{B}(\pi(p),r)$ *for any $p \in \mathcal{X}, r > 0$ (see Definition* 3.7*).*

F. Gluing and doubling

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

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

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

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

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

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

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

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

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

- For any $x \in \mathcal{W}$ we have $x \in \mathcal{V}$ or $x' \in \mathcal{V}$, and
$$|x' - y|_{\mathcal{W}} = |x - y'|_{\mathcal{W}} = \inf_{a \in A}\{|x - a|_{\mathcal{V}} + |a - y|_{\mathcal{V}}\}$$
for any $x, y \in \mathcal{V}$.

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

The following proposition follows directly from the definitions.

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

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

(b) *If $\mathcal{V}$ is proper, then so is $\mathcal{W}$. In this case, for any $x, y \in \mathcal{V}$ there is $a \in A$ such that*
$$|x - a|_{\mathcal{V}} + |a - y|_{\mathcal{V}} = |x - y'|_{\mathcal{W}}.$$

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

G. Kuratowsky embedding

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

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

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

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

Chapter 3

Maps and functions

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

A. Submaps

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

A submap $\Phi\colon \mathcal{X} \multimap \mathcal{Y}$ is *continuous* if the inverse image of any open set is open. Since $\operatorname{Dom}\Phi = \Phi^{-1}(\mathcal{Y})$, the domain $\operatorname{Dom}\Phi$ of a continuous submap is open. The same holds for upper and lower semicontinuous functions $f\colon \mathcal{X} \multimap \mathbb{R}$ since they are continuous functions for a special topology on $\mathbb{R}$.

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

B. Lipschitz conditions

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

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

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

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

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

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

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

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

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

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

for any $R \geqslant 0$ and $x \in \mathcal{X}$. A dual version of this property is considered in the following definition.

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

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

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

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

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

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

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

In other words, ℓ-co-Lipschitz maps can be considered as a quantitative version of open maps. For that reason they are also called ℓ-open [**44**]. Also, be aware that some authors refer to our ℓ-co-Lipschitz maps as $\frac{1}{\ell}$-co-Lipschitz.

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

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

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

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

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

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

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

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

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

Consider the function

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

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

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

C. Isometries and submetries

3.6. Isometry. *Let $\mathcal{X}$ and $\mathcal{Y}$ be metric spaces, and let $\Phi : \mathcal{X} \to \mathcal{Y}$ be a map.*

(a) *The map Φ is* distance-preserving *if*

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

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

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

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

3.7. Submetry. *A map $\sigma : \mathcal{X} \to \mathcal{Y}$ between the metric spaces $\mathcal{X}$ and $\mathcal{Y}$ is called a* submetry *if*

$$\sigma(\mathrm{B}(p, r)_{\mathcal{X}}) = \mathrm{B}(\sigma(p), r)_{\mathcal{Y}}$$

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

Note $\sigma : \mathcal{X} \to \mathcal{Y}$ is a submetry if it is 1-Lipschitz and 1-co-Lipschitz.

Note also that any submetry is an onto map.

The main source of examples of submetries comes from isometric group actions. Namely, assume $\mathcal{X}$ is a metric space and G is a subgroup of isometries of $\mathcal{X}$. Denote by $[x] = G\cdot x$ the G-orbit of $x \in \mathcal{X}$ and $\mathcal{X}/G$ the set of all G-orbits; let us equip it with the pseudometric defined by

$$|[x] - [y]|_{\mathcal{X}/G} = \inf\{\, |g\cdot x - h\cdot y|_{\mathcal{X}} : g, h \in G \,\}.$$

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

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

Proof. We need to show that the map $x \mapsto [x] = G\cdot x$ is 1-Lipschitz and 1-co-Lipschitz. The co-Lipschitz part follows directly from the definitions of Hausdorff distance and co-Lipschitz maps.

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

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

3.9. Proposition. *Let $\mathcal{X}$ be a length space, and let $\sigma\colon \mathcal{X} \to \mathcal{Y}$ be a submetry. Then $\mathcal{Y}$ is a length space.*

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

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

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

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

The curve $\sigma\circ\gamma$ joins x to y. Since σ is 1-Lipschitz and by the above,

$$\begin{aligned}\operatorname{length}\sigma\circ\gamma &\leqslant \operatorname{length}\gamma \\ &\leqslant |x-y|_{\mathcal{Y}} + \varepsilon.\end{aligned}$$

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

D. Speed of curves

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

The following theorem follows from [**37**, 2.7].

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

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

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

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

where $\int$ denotes the Lebesgue integral.

A curve $\alpha\colon \mathbb{I} \to \mathcal{X}$ is *unit-speed* if

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

for any subinterval $[a, b] \subset \mathbb{I}$. If α is Lipschitz, then, according to Theorem 3.10, this is equivalent to

$$\operatorname{speed} \alpha \overset{\text{a.e.}}{=\!=} 1.$$

The following generalization of the standard Rademacher theorem on differentiability almost everywhere (a.e.) of Lipschitz maps between smooth manifolds [**37**, 5.5.2] was proved by Bernd Kirchheim [**95**].

The conclusion of the standard Rademacher theorem does not make sense for maps to a metric space since the target might have no linear structure. But the theorem does not hold even if we assume that the target is a Banach space. For example, consider the map $[0, 1] \to L^1[0, 1]$, defined by $x \mapsto \chi_{[0,x]}$, where χ_A denotes the characteristic function of A. This map is distance-preserving and in particular Lipschitz, but its differential is undefined at any point.

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

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

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

E. Convex real-to-real functions

In this section we will discuss generalized solutions of the differential inequalities

(1) $$y'' + \kappa \cdot y \geqslant \lambda \quad \text{and, respectively,} \quad y'' + \kappa \cdot y \leqslant \lambda$$

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

Inequalities (1) are understood in the sense of distributions. That is, for any smooth function ϕ with compact support $\operatorname{Supp}\phi \subset \operatorname{Dom} y$, the following inequality should be satisfied:

$$\int\limits_{\operatorname{Dom} y} [y(t)\cdot\phi''(t) + \kappa\cdot y(t)\cdot\phi(t) - \lambda]\cdot \boldsymbol{d}t \geqslant 0, \tag{2}$$

$$\text{respectively} \quad \leqslant 0.$$

The integral is understood in the sense of Lebesgue; in particular inequality (2) makes sense for any Borel-measurable subfunction y. The proofs of the following propositions are straightforward.

3.12. Proposition. *Let $\mathbb{I} \subset \mathbb{R}$ be an open interval, and let $y_n\colon \mathbb{I} \to \mathbb{R}$ be a sequence of solutions of one of the inequalities in* (1)*. Assume $y_n(t) \to y_\infty(t)$ as $n \to \infty$ for any $t \in \mathbb{I}$. Then y_∞ is a solution of the same inequality in* (1)*.*

Assume y is a solution of one of the inequalities in (1). For $t_0 \in \operatorname{Dom} y$, let us define the *right* (*respectively, left*) *derivative* $y^+(t_0)$ ($y^-(t_0)$) at t_0 by

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

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

3.13. Proposition. *Let $\mathbb{I} \subset \mathbb{R}$ be an open interval, and let $y\colon \mathbb{I} \to \mathbb{R}$ be a solution of an inequality in* (1)*. Then y is locally Lipschitz; its right and left derivatives $y^+(t_0)$ and $y^-(t_0)$ are defined for any $t_0 \in \mathbb{I}$. Moreover*

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

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

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

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

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

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

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

such that

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

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

Further, the following property holds:

(d) *Suppose $y'' \leqslant \lambda - \kappa \cdot y$. Let $t_0 \in \mathbb{I}$, and let $\bar{y}$ be a solution of the ordinary differential equation $\bar{y}'' = \lambda - \kappa \cdot \bar{y}$ such that $\bar{y}(t_0) = y(t_0)$ and $y^+(t_0) \leqslant \leqslant \bar{y}^+(t_0) \leqslant -y^-(t_0)$. (Note that such a $\bar{y}$ is unique if y is differentiable at t_0.)*

Then $\bar{y} \geqslant y$ for all $t \in [t_0 - \varpi^\kappa, t_0 + \varpi^\kappa] \cap \mathbb{I}$; that is, $\bar{y}$ is a barrier of y at t_0. (Similarly, by reversing inequalities, for $y'' \geqslant \lambda - \kappa \cdot y$.)

Note that Theorem 3.14 implies that y satisfies $y'' \geqslant \lambda$ ($y'' \leqslant \lambda$) on an interval $\mathbb{I} \subset \mathbb{R}$ if and only if $y(t) - \frac{\lambda}{2} \cdot t^2$ is convex (respectively, concave) on $\mathbb{I}$.

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

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

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

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

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

Passing to the limit, we get

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

for $t \geqslant t_0$, and

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

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

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

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

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

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

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

Proof. Set $\hat{f}_n(t) = f_n(t) - \frac{\lambda}{2} \cdot t^2$ and $\hat{f}_\infty(t) = f_\infty(t) - \frac{\lambda}{2} \cdot t^2$. Note that the $\hat{f}_n$ are concave and $\hat{f}_n \to \hat{f}_\infty$ pointwise. Thus the theorem follows from 3.15, the Two-shoulder lemma. □

F. Convex functions on a metric space

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

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

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

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

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

in the domain $\{\, t : y(t) < \infty \,\}$; see the definition in Section 3.E.

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

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

if the subfunction $h = -f$ satisfies

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

Functions satisfying the inequalities

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

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

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

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

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

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

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

The proof is straightforward.

Chapter 4

Ultralimits

Here we introduce ultralimits of sequences of points, metric spaces, and functions. Our presentation is based on [**97**].

Ultralimits are closely related to Gromov–Hausdorff limits. We use them only as a canonical way to pass to convergent subsequences. We could avoid using them at the cost of saying “pass to a convergent subsequence” too many times. Doing this might be cumbersome and it obscures ideas of the proof; see for example the proof of the globalization theorem for general CBB spaces in Section 8.F. Also the use of ultralimits is convenient when dealing with CAT spaces due to the lack of compactness results.

A. Ultrafilters

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

Measure-theoretic definition. Recall that $\mathbb{N}$ denotes the set of natural numbers, $\mathbb{N} = \{1, 2, \dots\}$.

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

(a) $\omega(S) = 0$ or 1 for any subset $S \subset \mathbb{N}$. An ultrafilter ω is called *nonprincipal* if in addition

(b) $\omega(F) = 0$ for any finite subset $F \subset \mathbb{N}$. A nonprincipal ultrafilter ω is called *selective* if in addition

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

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

4.2. Advanced exercise. Let ω be an ultrafilter, and let $f : \mathbb{N} \to \mathbb{N}$. Suppose that $\omega(S) \leqslant \omega(f^{-1}(S))$ for any set $S \subset \mathbb{N}$. Show that $f(n) = n$ for ω-almost all $n \in \mathbb{N}$.

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

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

Setting

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

makes these two definitions equivalent.

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

The existence of a selective ultrafilter follows from the continuum hypothesis; this was proved by Walter Rudin [**144**].

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

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

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

B. Ultralimits of points

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

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

In this case, we write

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

Also, if $\mathcal{X} = \mathbb{R}$ we write $\lim_{n\to\omega} x_n = \pm\infty$ if

$$\omega\{\, n \in \mathbb{N} : \pm x_n > L \,\} = 1$$

for any $L \in \mathbb{R}$.

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

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

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

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

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

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

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

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

Now assume that ω is selective. Consider the sets

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

4.6. Exercise. Recall that ℓ^∞ denotes the space of bounded sequences of real numbers. Show that there is a linear functional $L\colon \ell^\infty \to \mathbb{R}$ such that for any sequence $\boldsymbol{s} = (s_1, s_2, \dots) \in S$ the image $L(\boldsymbol{s})$ is a partial limit of $s_1, s_2, \dots$.

4.7. Exercise. Suppose that $f\colon \mathbb{N} \to \mathbb{N}$ is a map such that

$$\lim_{n\to\omega} x_n = \lim_{n\to\omega} x_{f(n)}$$

for any bounded sequence x_n of real numbers. Show that $f(n) = n$ for ω-almost all $n \in \mathbb{N}$.

C. Ultralimits of spaces

Fix a selective ultrafilter ω on the set of natural numbers.

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

$$|(x_n) - (y_n)| = \lim_{n\to\omega} |x_n - y_n|. \tag{1}$$

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

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

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

and the distance is defined as in (1).

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

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

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

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

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

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

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

$$\mathbb{N} = S_1 \supset S_2 \supset \cdots,$$

such that

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

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

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

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

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

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

By Observation 4.8, $\mathcal{X}_\omega$ is complete. Applying Lemma 2.13, we obtain the statement. ☐

A geodesic space $\mathcal{T}$ is called a *metric tree* if any pair of points in $\mathcal{T}$ are connected by a unique geodesic, and the union of any two geodesics $[xy]_\mathcal{T}$, and $[yz]_\mathcal{T}$ contains the geodesic $[xz]_\mathcal{T}$. The latter means that any triangle in $\mathcal{T}$ is a tripod; that is, for any three points x, y, and z there is a point p such that

$$[xy] \cup [yz] \cup [zx] = [px] \cup [py] \cup [pz].$$

4.10. Exercise. Let $\mathcal{T}$ be a metric component of the ultralimit of $\mathbb{M}^2(n)$ as $n \to \to \omega$.

(a) Show that $\mathcal{T}$ is a complete metric tree.

(b) Show that $\mathcal{T}$ is homogeneous; that is, given two points $s, t \in \mathcal{T}$ there is an isometry of $\mathcal{T}$ that maps s to t.

(c) Show that $\mathcal{T}$ has *continuum degree* at any point; that is, for any point $t \in \mathcal{T}$ the set of connected components of the complement $\mathcal{T} \setminus \{t\}$ has cardinality continuum.

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

By Theorem 5.16, there is a distance-preserving map $\iota\colon \mathcal{X} \hookrightarrow \mathcal{X}^\omega$, where $\iota(y)$ is the equivalence class of the constant sequence $y_n = y$.

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

According to Theorem 5.16, if $\mathcal{X}$ is compact, then $\iota(\mathcal{X}) = \mathcal{X}^\omega$; in particular, $\mathcal{X}^\omega$ is isometric to $\mathcal{X}$. If $\mathcal{X}$ is proper, then $\iota(\mathcal{X})$ forms a metric component of $\mathcal{X}^\omega$.

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

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

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

Note that $|x - z_n|_{\mathcal{X}} \to \frac{1}{2}\cdot|x - y|_{\mathcal{X}}$ and $|y - z_n|_{\mathcal{X}} \to \frac{1}{2}\cdot|x - y|_{\mathcal{X}}$ as $n \to \omega$. In particular, for any $\varepsilon > 0$, the point z_n is an ε-midpoint between x and y for ω-almost all n. It remains to apply Lemma 2.13.

The "if" part follows from Observation 4.9. □

Note that the proof above together with Lemma 4.5 imply the following:

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

4.13. Exercise. Let $\mathcal{X}$ be a countable set with discrete metric; that is $|x-y|_{\mathcal{X}} = 1$ if $x \ne y$. Show that

(a) $\mathcal{X}^\omega$ is not isometric to $\mathcal{X}$.

(b) $\mathcal{X}^\omega$ is isometric to $(\mathcal{X}^\omega)^\omega$.

4.14. Exercise. Given a nonprincipal ultrafilter ω, construct an ultrafilter ω_1 such that

$$\mathcal{X}^{\omega_1} \stackrel{\text{iso}}{=\!=} (\mathcal{X}^\omega)^\omega$$

for any metric space $\mathcal{X}$.

4.15. Exercise. Construct a proper metric space $\mathcal{X}$ such that $\mathcal{X}^\omega$ is not proper; that is, there is a point $p \in \mathcal{X}^\omega$ and $R < \infty$ such that the closed ball $\overline{\mathrm{B}}[p, R]_{\mathcal{X}^\omega}$ is not compact.

D. Ultralimits of sets

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

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

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

Applying Observation 4.8 to the sequence of complements $\mathcal{X}_n \setminus \Omega_n$, we see that Ω_ω is open for any sequence Ω_n.

This definition can be applied to arbitrary sequences of sets, but we will apply it only for sequences of open sets.

E. Ultralimits of functions

Recall that a family of submaps (see Section 3.A) between metric spaces $\{f_\alpha \colon \mathcal{X} \multimap \mathcal{Y}\}_{\alpha\in\mathcal{A}}$ is called *equicontinuous* if for any $\varepsilon > 0$ there is $\delta > 0$ such that for any $p, q \in \mathcal{X}$ with $|p-q| < \delta$ and any $\alpha \in \mathcal{A}$, we have $|f_\alpha(p)-f_\alpha(q)| < \varepsilon$.

Let $f_n : \mathcal{X}_n \circ\!\!\longrightarrow \mathbb{R}$ be a sequence of subfunctions.

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

Assume there is a subfunction $f_\omega : \mathcal{X}_\omega \circ\!\!\longrightarrow \mathbb{R}$ that satisfies the following conditions:

(1) $\operatorname{Dom} f_\omega = \Omega_\omega$,

(2) if $x_n \to x_\omega \in \Omega_\omega$ for a sequence of points $x_n \in \mathcal{X}_n$, then $f_n(x_n) \to f_\omega(x_\omega)$ as $n \to \omega$.

In this case, the subfunction $f_\omega : \mathcal{X}_\omega \to \mathbb{R}$ is said to be the ω-limit of $f_n : \mathcal{X}_n \to \mathbb{R}$.

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

4.16. Lemma. *Let $\mathcal{X}_n$ be a sequence of metric spaces, and let $f_n : \mathcal{X}_n \circ\!\!\longrightarrow \mathbb{R}$ be a sequence of subfunctions.*

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

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

The proof is straightforward.

4.17. Exercise. Construct a sequence of compact length spaces $\mathcal{X}_n$, with a converging sequence of ℓ-Lipschitz concave (see Definition 3.17) functions $f_n : \mathcal{X}_n \to \mathbb{R}$ such that the ω-limit $\mathcal{X}_\omega$ of $\mathcal{X}_n$ is compact and the ω-limit $f_\omega : \mathcal{X}_\omega \to \mathbb{R}$ of f_n is not concave.

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

$$f^\omega : \mathcal{X} \circ\!\!\longrightarrow \mathbb{R}, \quad f^\omega(x_\omega) = \lim_{n\to\omega} f(x_n).$$

Evidently, if f^ω is defined, then f is continuous.

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