Geometry & Groups

Stephan Tillmann

(Version 8.1.2012)

2 Metric spaces

These notes collect and augment facts about metric spaces and manifolds that can be found in Bonahon [1]. Good references for this handout are Folland [19], Munkres [28] and Patty [30]. Some of the material about metric spaces should be familiar from analysis or multivariable calculus if *metric space* is replaced by $\mathbb{R}$ or $\mathbb{R}^n$. Manifolds will usually arise in a very concrete way in this course: either from gluing together pieces of $\mathbb{R}^n$ or as the quotient of a group action on some existing manifold. Through either of these constructions, they will have a *quotient metric* ([1]§§4.2 and 7.2), which we use to analyse them.

Not all of the material in these notes will be used or referenced during the course: various cheerful but irrelevant facts have been added to give a small glimpse of the rich theory of metric spaces and manifolds. These peripheral results are discussed in the hope that they help clarify the core material or put it in context. If a whole section consists of such material, it is marked with an asterisk.

2.1 Metric spaces

(Related to [1] §1.3.)

Definition 2.1 *A* metric space *is a pair (X,d) consisting of a set X and a function $d : X \times X \to \mathbb{R}$ such that*

(1) $d(x,y) \geq 0$ *and* $d(x,x) = 0$ *for all* $x,y \in X$;

(2) $d(x,y) = 0$ *only if* $x = y$;

(3) $d(x,y) = d(y,x)$;

(4) $d(x,y) \leq d(x,z) + d(z,y)$.

The function d is called a metric *on X. If a function $d : X \times X \to \mathbb{R}$ only satisfies (1), (3) and (4), then it is called a* semi-metric *on X.*

The classic example of a metric space is the set of all real numbers with the euclidean distance function:

$$d(x,y) = |x-y|.$$

A more exotic example would be any random set X and $d(x,y) = 1$ if $x \neq y$ and $d(x,y) = 0$ otherwise. This is called the *discrete metric* on X.

A *subspace* of a metric space (X,d) is a subset $Y \subseteq X$ with the *induced metric* $d|_{Y\times Y} \colon Y \times Y \to \mathbb{R}$. Then $(Y, d|_{Y\times Y})$ is a metric space in its own right. This metric is usually put on subsets of a metric space.

A different metric that can often be put on a subset is the *path metric* (also called *inner metric*); see §2.13.

2.2 Isometry

(Related to [1] §1.4)

Given the metric spaces (X,d) and (X',d'), a map $f\colon X \to X'$ is *distance preserving* if

$$d(x,y) = d'(f(x), f(y))$$

for all $x,y \in X$. It follows from the definition of a metric that all distance preserving maps are injective. A distance preserving bijection is called an *isometry*. If an isometry from (X,d) to (X',d') exists, then the spaces are called *isometric*. The inverse of an isometry is clearly an isometry, and so are the identity map and the composition of any two isometries. The group of all isometries from a metric space (X,d) to itself is denoted $\mathrm{Isom}(X,d)$ or, if d is understood, $\mathrm{Isom}(X)$.

2.3 Diameter

Definition 2.2 (Diameter) *Suppose $\emptyset \neq A \subseteq X$. Then*

$$\mathrm{diam}(A) = \sup\{d(x,y) \mid x,y,\in A\} \in [0,\infty].$$

The set A is bounded *if* $\mathrm{diam}(A) < \infty$.

Note that $A \subseteq B$ implies $\mathrm{diam}(A) \leq \mathrm{diam}(B)$. It can happen that (X,d) is bounded; e.g. X with the discrete metric has diameter equal to 1.

2.4 Convergence and completeness

(Related to [1] §1.3 and §6.4.)

Convergence in metric spaces is defined just as in calculus.

Definition 2.3 (Convergence and limit) *A sequence $(x_n)_{n\in\mathbb{N}}$ of points in a metric space (X,d) converges to a limit $x_\infty \in X$ if*

$$\forall \varepsilon > 0\ \exists N \in \mathbb{N} \text{ such that } n > N \text{ implies } d(x_n, x_\infty) < \varepsilon.$$

This is often written as $x_n \to x_\infty$. We say that $(x_n)_{n\in\mathbb{N}}$ is convergent if it converges to some point in X.

If a limit x_∞ exists, it is unique, and so x_∞ is called *the* limit of the sequence.

For instance, a sequence in $\mathbb{R}^n$ with the euclidean metric converges if and only if it converges componentwise.

Definition 2.4 (Cauchy sequence) *A sequence $(x_n)_{n\in\mathbb{N}}$ of points in a metric space (X,d) is Cauchy if*

$$\forall \varepsilon > 0\ \exists N \in \mathbb{N} \text{ such that } n,m > N \text{ implies } d(x_n, x_m) < \varepsilon.$$

Lemma 2.5 *A convergent sequence is Cauchy.*

Definition 2.6 (Complete) *A subset A of a metric space (X,d) is complete if every Cauchy sequence in A is convergent and has its limit in A.*

Theorem 2.7 (Completion) *Every metric space can be realized as a subspace of a complete metric space.*

For instance, $\mathbb{Q}$ with the usual euclidean metric is *not* complete, but it is a subspace of $\mathbb{R}$, which is. This example also gives an idea of the proof: One carefully adds limits for all the non-convergent Cauchy sequences to the incomplete space in order to get *the completion.* A completion $\overline{X}$ of X is a complete space that contains an isometric copy of X that is dense in $\overline{X}$; it turns out to be unique up to isometry.

2.5 Open and closed sets in a metric space

(Related to [1] §1.3.)

Definition 2.8 (Open and closed in metric space) *If (X,d) is a metric space, then the* (open) ball *with centre $x \in X$ and radius $r > 0$ is:*

$$B_d(x,r) = \{y \in X : d(x,y) < r\}.$$

We then say that $V \subseteq X$ is open if for all $x \in V$ there exists $r = r(x) > 0$ with

$$B_d(x,r) \subseteq V.$$

A subset of X is closed if its complement is open.

It follows from the definition that X and the empty set are both open and hence both closed. One can check that an open ball in X is indeed an open set in X. Moreover, arbitrary unions of open sets are open; finite intersections of open sets are open; arbitrary intersections of closed sets are closed; and finite unions of closed sets are closed.

If X has the discrete metric, then every subset of X is both open and closed. In general, there may be some subsets of X that are neither open nor closed; for example the interval $[0,1)$ is neither open nor closed in $(\mathbb{R}, d_{euc})$.

Lemma 2.9 (Closure = sequential closure in metric spaces) *A subset A of a metric space X is closed if and only if whenever $(x_n)_{n\in\mathbf{N}}$ is a convergent sequence of points in A, the limit is also in A.*

Corollary 2.10 *A closed subset of a complete metric space is complete, and a complete subset of an arbitrary metric space is closed.*

Here is another cheerful fact that may be known from calculus:

Proposition 2.11 (The nesting principle) *Let (X,d) be a complete metric space and $A_0 \supseteq A_1 \supseteq A_2 \supseteq \ldots$ be a nested sequence of closed, non-empty subsets with* $\operatorname{diam}(A_n) \to 0$. *Then $\bigcap A_n$ is non-empty and consists of a single point.*

2.6 Topological spaces

The notions of *open* and *closed* sets in a metric space provide a nice framework for the study of general properties of the space. It is useful to move to the next step of abstraction straight away: the concept of a *topological space*, where one has the notions of open and closed sets even if no metric is present.

For a set X, let $\mathscr{P}(X)$ denote the set of all subsets of X.

Definition 2.12 *A topological space $(X,\mathscr{O})$ consists of a set X and a set $\mathscr{O} \subseteq \mathscr{P}(X)$, such that*

(1) $\emptyset, X \in \mathscr{O}$,

(2) *$U, V \in \mathscr{O}$ implies $U \cap V \in \mathscr{O}$,*

(3) *if $U_\alpha \in \mathscr{O}$ for each α in an arbitrary index set A, then $\bigcup_{\alpha\in A} U_\alpha \in \mathscr{O}$.*

An element of $\mathscr{O}$ is called an open set *of $(X,\mathscr{O})$, and a subset of X is* closed *if its complement is open.*

The set of all open sets in a metric space (X,d) (as per §2.5) defines a topology on X, and this is called the *metric topology*. Whence one can refer to open sets in a metric space without ambiguity, and I'll always assume that a metric space is given the metric topology.

Another example is given by the *discrete topology* on a set X. Here $\mathscr{O} = \mathscr{P}(X)$, so every subset of X is both open and closed. This is the topology that arises from the discrete metric on X. This is the *finest* topology that exists on X.

In contrast, the *coarsest* topology on the set X is $\mathscr{O} = \{\emptyset, X\}$. If X contains at least two distinct elements, then this topology cannot arise from a metric on X. This follows from the following definition and exercise.

Definition 2.13 (Hausdorff) *The topological space $(X,\mathscr{O})$ is* Hausdorff *if for any $x,y \in X$ with $x \neq y$, there there exist open sets U and V such that $x \in U$, $y \in V$ and $U \cap V = \emptyset$.*

Exercise 2.14 (Metric topology is Hausdorff) *If (X,d) is a metric space, then it is Hausdorff.*

A discussion of other *separation axioms* is given in §2.16. In many of the definitions below, it is useful to have metric spaces in mind at first reading. The definition of convergence in metric spaces generalises as follows:

Definition 2.15 (Convergence) *A sequence $(x_n)_{n\in\mathbb{N}}$ in a topological space $(X,\mathscr{O})$ converges to a point $x \in X$ if for each open neighbourhood U of x there exists $N \in \mathbb{N}$ such that $x_n \in U$ for all $n \geq N$.*

In analogy to the restriction of a metric on a subspace, we can also restrict the topology to a subspace:

Definition 2.16 (Subspace topology) *Let Y be a subset of the topological space $(X,\mathscr{O})$. The* subspace topology *on Y determined by $\mathscr{O}$ is $\mathscr{O}_Y = \{Y \cap U \mid U \in \mathscr{O}\}$.*

For a subset of a metric space, it is clear that the induced metric gives the subspace topology. For example, $\mathbb{Z} \subset \mathbb{R}$ inherits the discrete topology from the usual topology on $\mathbb{R}$, since for every integer $n \in \mathbb{Z}$, we have $\{n\} = \mathbb{Z} \cap (n - \frac{1}{2}, n + \frac{1}{2}) \in \mathscr{O}_{\mathbb{Z}}$.

Suppose we have a topological space $(X,\mathscr{O})$, and a surjective function $f\colon X \to Y$, where Y is a set. We'd now like to transfer the topology from X to Y in a natural way. Since f is a surjection, the sets $f^{-1}(y)$ form a partition of X, since they are non-empty and f is a function. We therefore obtain an equivalence relation on X by declaring $x \sim y$ if and only if $f(x) = f(y)$, and the set of equivalence classes can be identified with Y. The topology is transferred as follows:

Definition 2.17 (Quotient topology) *Suppose $\sim$ is an equivalence relation on the topological space $(X,\mathscr{O})$. Define the* quotient topology *on the set of equivalence classes by decreeing that a set of classes is open if and only if the union of their inverse images in X is open.*

2.7 Interior, closure and boundary

A *neighbourhood* of a point $x \in X$ is any set $U \subseteq X$ such that there is an open set V satisfying $x \in V \subseteq U$. Note that a neighbourhood is not necessarily open nor closed. For example, if (X,d) is a metric space, then $B_d(x,r)$ is a perfectly fine neighbourhood of x, and so is X (but the latter is often not so useful).

Definition 2.18 (Interior, closure and boundary) *Given a topological space $(X,\mathscr{O})$ and a subset $A \subseteq X$, define*

$$\operatorname{int}(A) = \{x \in A \mid \text{there is a neighbourhood } U \text{ of } x \text{ such that } x \in U \subset A\} \text{ (interior of } A\text{)}$$

$$\operatorname{cl}(A) = \{x \in X \mid \text{for each neighbourhood } U \text{ of } x, \text{ we have } U \cap A \neq \emptyset\} \text{ (closure of } A\text{)}$$

$$\partial A = \{x \in X \mid \text{for each neighbourhood } U \text{ of } x, \text{ we have } U \cap A \neq \emptyset \neq U \cap (X \setminus A)\} \text{ (boundary of } A\text{)}$$

Exercise 2.19 *The interior of A is the largest open set contained in A, the closure of A is the smallest closed set containing A, and $\partial A = \mathrm{cl}(A) \setminus \mathrm{int}(A)$. Moreover, A is open if and only if $A \cap \partial A = \emptyset$; A is closed if and only if $\partial A \subseteq A$.*

It should be noted that interior, closure and boundary depend on the ambient metric space X. For instance, if $X = [0,1)$ with the induced metric from $(\mathbb{R}, d_{euc})$, then $(0,1) \subset X$ has $\partial(0,1) = \{0\}$ and $\mathrm{cl}(0,1) = [0,1)$. If there is risk of confusion, one writes int_X, cl_X and ∂_X to emphasise the ambient space.

Definition 2.20 (Dense) *The subset A of X is* dense *in X if $\mathrm{cl}(A) = X$.*

For instance, the rationals are dense in the reals.

2.8 Continuous functions

(Related to [1] §1.3.)

Definition 2.21 (Continuous function) *For $(X, \mathscr{O})$ and $(X', \mathscr{O}')$ two topological spaces, we say that a function*

$$f : X \to X'$$

is continuous if $f^{-1}(U)$ is open in X for any open $U \subseteq X'$.

For instance, if (X,d) is $\mathbb{R}^2$ with the euclidean metric and (X',d') is $\mathbb{R}^2$ with the discrete metric, then the identity map on $\mathbb{R}^2$ is a continuous map $X' \to X$ but not a continuous map $X \to X'$.

It follows from the definition of the quotient topology, that if $\sim$ is an equivalence relation on the topological space $(X, \mathscr{O})$, and $X/\sim$ is given the quotient topology, then the natural map $X \to X/\sim$ is continuous.

The connection between the above definition and the customary notion for $\mathbb{R}$ is made by the following lemma:

Lemma 2.22 *A function $f : X \to X'$ between the metric space (X,d) and the metric space (X',d') is continuous if and only if for all $x \in X$ and $\delta > 0$ there exists $\varepsilon > 0$ such that*

$$\text{for all } y \in X,\ d(x,y) < \varepsilon \text{ implies } d'(f(x), f(y)) < \delta.$$

This directly shows that an isometry is continuous (simply choose $\delta = \varepsilon$ in the above lemma).

There also is a notion of *continuity at a point*:

Definition 2.23 (Continuous at point) *For $(X, \mathscr{O})$ and $(X', \mathscr{O}')$ two topological spaces, we say that a function*

$$f : X \to X'$$

is continuous at $x \in X$ if $f^{-1}(U)$ is a neighbourhood of x in X for any neighbourhood $U \subseteq X'$ of $f(x)$.

It is important for the above definition that a neighbourhood is not necessarily open, but only contains an open set containing the point in question. One can check that f is continuous if and only if it is continuous at every point, so the two definitions work well in tandem.

Once a metric comes into play, some continuous maps are nicer than others; see the example of a non-rectifiable path in §2.13. The following notion (which is irrelevant for this course) may be familiar from analysis:

Definition 2.24 (Uniform continuity) *For (X,d) and (X',d') two metric spaces, we say that the function $f : X \to X'$ is uniformly continuous if for all $\delta > 0$ there exists $\varepsilon > 0$ such that*

$$\text{for all } x, y \in X,\ d(x,y) < \varepsilon \text{ implies } d'(f(x), f(y)) < \delta.$$

It follows from Lemma 2.22 that a uniformly continuous function is continuous; it has the added advantage that δ only depends on ε but not on the specific point of X.

2.9 Contractions*

Definition 2.25 (Contraction) *For (X,d) and (X',d') two metric spaces, we say that the function $f : X \to X'$ is a* contraction *with* contraction constant $\lambda \in [0,1)$ *for all $x, y \in X$*

$$d'(f(x), f(y)) \leq \lambda d(x,y).$$

A contraction is continuous—just choose $\delta = \varepsilon\lambda$ in Lemma 2.22. A nice fact about contractions is the *contraction mapping principle* aka *Banach Fixed Point Theorem*:

Theorem 2.26 *Let $f\colon X \to X$ be a contraction of a complete metric space (X,d). Then f has a unique fixed point; i.e. there is a unique $p \in X$ such that $f(p) = p$.*

Moreover, for any $x \in X$ define $x_0 = x$ and $x_{n+1} = f(x_n)$ for $n \geq 0$. Then $x_n \to p$ and $d(x,p) \leq \dfrac{d(x,f(x))}{1-\lambda}$.

Uniqueness is easy to show. The idea of the existence proof is in the second statement: first show that the sequence $(x_n)_{n\in\mathbb{N}}$ arising from iterated application of f is Cauchy, and then use the continuity of f to show that its limit (X is complete!) is a fixed point.

The contraction mapping principle is rather fundamental. For instance, using it one can prove local existence and uniqueness of solutions to ordinary differential equations (Picard's theorem) and the inverse and implicit function theorems. (It is a fun exercise to show that the inverse and implicit function theorems follow from each other. However, proving either of them with the contraction mapping principle is quite a bit harder).

2.10 Homeomorphisms

(Related to [1] §5.1.)

Definition 2.27 (Homeomorphism) *For $(X,\mathscr{O})$ and $(X',\mathscr{O}')$ two topological spaces, we say that the function $f : X \to X'$ is a* homeomorphism *if it is bijective, continuous and its inverse is also continuous.*

Exercise 2.28 *Show that the map $[0,1) \to S^1 = \{z \in \mathbb{C} \mid |z| = 1\}$ defined by $t \mapsto e^{2\pi it}$ is continuous and injective, but that its inverse is not continuous. (Here both the domain and the range are given the induced metrics from the euclidean metrics on $\mathbb{R}$ and $\mathbb{C}$ respectively.)*

Exercise 2.29 *Show that $B_{d_{euc}}(x,r)$ is homeomorphic to $\mathbb{R}^n$ for each $x \in \mathbb{R}^n$ and each $r > 0$. (Here, d_{euc} is the standard Euclidean metric).*

Homeomorphism is an equivalence relation, and a *topological property* of a topological space is one that is not changed by homeomorphisms. We will encounter various basic topological properties such as connectedness (§2.12), compactness (§2.15) and metrizability (§2.16).

Isometric metric spaces are homeomorphic, but the converse is not true:

Exercise 2.30

(1) *Two metrics d and d' on the set X induce the same topology on X if and only if*

$$Id_X\colon\ (X,d) \to (X,d')$$

is a homeomorphism.

(2) *Give an example of two homeomorphic metric spaces that are not isometric.*

We are sometimes only interested in spaces up to homeomorphism, and in this case, a fixed metric is a useful vehicle for the study of continuous maps on the space.

2.11 The extended reals*

The *extended reals* is the set $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, \infty\}$. The following conventions are used for arithmetic with $\pm\infty$ given $a \in \mathbb{R}$:

$$\begin{aligned} a \pm \infty &= \pm\infty, & a\,(\pm\infty) &= \pm\infty \ (a > 0), \\ \infty + \infty &= \infty, & a\,(\pm\infty) &= \mp\infty \ (a < 0), \\ -\infty - \infty &= -\infty, & 0\,(\pm\infty) &= 0. \end{aligned}$$

All other operations (such as $\infty - \infty$) are left undefined. We also declare $-\infty < a < \infty$, and can therefore make sense of intervals such as $[-\infty, a)$ and $[a, \infty]$. We can turn $\overline{\mathbb{R}}$ into a metric space by defining

$$\begin{aligned} d(x,y) &= |\arctan(x) - \arctan(y)| \quad \text{if } x, y \in \mathbb{R}, \\ d(x,\infty) = d(\infty,x) &= \frac{\pi}{2} - \arctan(x) \quad \text{if } x \in \mathbb{R}, \\ d(x,-\infty) = d(-\infty,x) &= \arctan(x) + \frac{\pi}{2} \quad \text{if } x \in \mathbb{R}, \\ d(-\infty,\infty) &= \pi, \\ d(\pm\infty,\pm\infty) &= 0. \end{aligned}$$

Then $(\overline{\mathbb{R}}, d)$ is complete and isometric with $[-\frac{\pi}{2}, \frac{\pi}{2}]$.

2.12 Connectedness

There are various degrees of connectedness that one can consider. Here are two of them.

Definition 2.31 (Connected) *A topological space $(X, \mathscr{O})$ is* connected *if it cannot be written in the form $X = U \cup V$, where both U and V are non-empty open sets and $U \cap V = \emptyset$.*

A space is *disconnected* if it is not connected. A space is *totally disconnected* if the only connected subsets are the empty set and the sets consisting of a single point. An example is any set with the discrete topology/metric.

Theorem 2.32 (Continuous image of connected is connected) *Suppose $(X, \mathscr{O})$ is connected and $f\colon X \to X'$ is a continuous surjection onto an arbitrary topological space. Then $(X', \mathscr{O}')$ is also connected.*

One can show that any connected subset of $(\mathbb{R}, d_{euc})$ is an interval. A consequence of the above theorem is then the:

Corollary 2.33 (Intermediate Value Theorem) *Suppose $(X, \mathscr{O})$ is connected and $f\colon X \to \mathbb{R}$ is continuous. Suppose $a, b \in X$ with $f(a) < f(b)$ and let $d \in R$ such that $f(a) < d < f(b)$. Then there is $c \in X$ such that $f(c) = d$.*

Definition 2.34 (Components) *Define an equivalence relation on $(X, \mathscr{O})$ by $x \sim y$ if and only if there is a connected subset $A \subseteq X$ such that $x \in A \ni y$. The equivalence classes are the* components *of X.*

One can show that the closure of a connected set is connected. Hence every component of X is closed. However, it is not necessarily open.

Throughout, $I = [0,1]$ denotes the unit interval and $[a,b] \subset \mathbb{R}$ a closed interval and (if thought of as metric or topological spaces) they are given the usual Euclidean metric and the induced topology unless stated otherwise.

Definition 2.35 (Path) *A continuous map $f\colon [a,b] \to X$ is a* path *or* curve *in $(X, \mathscr{O})$ from $f(a)$ to $f(b)$.*

Definition 2.36 (Path components) *Define an equivalence relation on $(X,\mathscr{O})$ by $x \sim y$ if and only if there is a path in X from x to y. The equivalence classes are the* path components *of X.*

It follows from the above that every path component is contained in a component of X; but the converse is not true in general.

Definition 2.37 (Path connected) *A topological space $(X,\mathscr{O})$ is* path connected *if for any two points $x,y \in X$ there is a path $f\colon I \to X$ with $f(0) = x$ and $f(1) = y$.*

Here are some nice standard applications using connectedness/disconnectedness:

(1) Any continuous function $f\colon [0,1] \to [0,1]$ has a fixed point. (Consider $g(x) = x - f(x)$.)

(2) $\mathbb{R}$ is not homeomorphic with $\mathbb{R}^n$ for any $n \geq 2$. (What happens when you remove a point?)

(3) $GL_n(\mathbb{R})$ is not connected. (Consider the determinant.)

2.13 The path metric*

(Related to [1] Exercises 1.10, 1.11 and 1.12.)

Definition 2.38 (Length of path) *Let (X,d) be a metric space and $\gamma\colon [a,b] \to X$ be a path. The length $\ell_d(\gamma) \in [0,\infty]$ is defined by*

$$\ell_d(\gamma) = \sup\{\sum_{k=1}^{n} d(\gamma(t_{k-1}),\gamma(t_k)) \mid a = t_0 < \ldots < t_n = b\},$$

where the supremum is taken over all partitions (no bound on n) of the interval $[a,b]$.

If $\ell_d(\gamma) < \infty$, then γ is called rectifiable.

An example of a path $\gamma\colon [0,1] \to [0,1]$ that is not rectifiable is given by:

$$\gamma(t) = t\cos\left(\frac{\pi}{t}\right) \text{ for } t \neq 0, \text{ and } \gamma(0) = 0.$$

Let $\gamma\colon [a,b] \to X$ be a path in a metric space (X,d). The following facts are implied by the definition:

(1) $\ell_d(\gamma) \geq d(\gamma(a),\gamma(b))$ and $\ell_d(\gamma) = 0$ if and only if γ is constant.

(2) If $\varphi\colon [a',b'] \to [a,b]$ is weakly monotonic, then $\ell_d(\gamma) = \ell_d(\gamma \circ \varphi)$.

(3) The reverse path $\overline{\gamma}\colon [a,b] \to X$ defined by $\overline{\gamma}(t) = (b+a-t)$ satisfies $\ell_d(\overline{\gamma}) = \ell_d(\gamma)$.

(4) Let $c \in (a,b)$ and write γ_1 for the restriction of γ to $[a,c]$ and γ_2 for the restriction of γ to $[c,b]$. Then $\ell_d(\gamma) = \ell_d(\gamma_1) + \ell_d(\gamma_2)$.

Definition 2.39 *(Path metric) Let (X,d) be a metric space, and $Y \subseteq X$ be given the induced metric. The* path metric *on Y is the map $\overline{d}\colon Y \times Y \to \overline{\mathbb{R}}$ defined by*

$$\overline{d}(x,y) = \inf \ell_d(\gamma),$$

where the infimum is taken over all rectifiable paths $\gamma\colon [0,1] \to Y$ with $\gamma(0) = x$ and $\gamma(1) = y$; if there is no such path, then $\overline{d}(x,y) = \infty$.

We have $\overline{d}(x,y) \geq d(x,y)$ for all $x,y \in Y$, and $\overline{d}$ indeed defines a (possibly $[0,\infty]$–valued) metric on Y. Notice that there may be no path of length $d(x,y)$ from x to y, even if $\overline{d}(x,y) = d(x,y)$ and $Y = X$.

2.14 Product metric and product topology*

(Related to [1] Exercise 1.6.)

Given metric spaces (X,d) and (X',d'), the following define metrics on $X \times X'$:

$$\big((x,y),(a,b)\big) \mapsto \max\{d(x,a), d'(y,b)\}$$
$$\big((x,y),(a,b)\big) \mapsto d(x,a) + d'(y,b)$$
$$\big((x,y),(a,b)\big) \mapsto \sqrt{(d(x,a))^2 + (d'(y,b))^2}$$

Exercise 2.40 *Show that the above metrics induce the same topology on $X \times X'$.*

The first of the above metrics is usually called the *product metric* on $X \times X'$. There also is a notion of a *product topology*, which we will only need for the product of two topological spaces (where it mimics the construction of the product metric).

Definition 2.41 (Product topology) *Let $(X,\mathscr{O})$ and $(X',\mathscr{O}')$ be two topological spaces. The product topology on $X \times X'$ is the subset of $\mathscr{P}(X \times X')$ consisting of every set that can be written as a union of sets of the form $U \times V$, where $U \in \mathscr{O}$ and $V \in \mathscr{O}'$.*

It is not difficult to check that the product topology satisfies the definition of a topology; the main thing to check is that the intersection of two sets of the form $U \times V$ can also be written in this form. Since $\emptyset \times \emptyset = \emptyset$, the product topology contains the empty set. It also contains $X \times X'$ by virtue of $X \in \mathscr{O}$ and $X' \in \mathscr{O}'$.

2.15 Compactness

(Related to [1] §6.4.)

Compactness is a fundamental property that can be viewed as generalising the essence of "finiteness".

Definition 2.42 (Cover) *Let $(X,\mathscr{O})$ be a topological space and $Y \subseteq X$. Suppose $\{U_\alpha \mid \alpha \in A\}$ is a collection of subsets of X with $Y \subseteq \bigcup_{\alpha\in A} U_\alpha$, then $\{U_\alpha \mid \alpha \in A\}$ is said to be a* cover *of Y.*

If each U_α is open (closed, compact, ...), then $\{U_\alpha \mid \alpha \in A\}$ is an open (closed, compact, ...) cover *of Y.*

If A is a finite set, then $\{U_\alpha \mid \alpha \in A\}$ is a finite cover *of Y.*

If $B \subseteq A$, and $Y \subseteq \bigcup_{\alpha\in B} U_\alpha$, then $\{U_\alpha \mid \alpha \in B\}$ is said to be a subcover *of $\{U_\alpha \mid \alpha \in A\}$.*

Definition 2.43 (Compact) *In a topological space $(X,\mathscr{O})$ we say that $Y \subseteq X$ is compact if every open cover of Y has a finite subcover; i.e. whenever $\{U_\alpha \mid \alpha \in A\}$ is a collection of open sets with $Y \subseteq \bigcup_{\alpha\in A} U_\alpha$, then there is some finite set $F \subseteq A$ with $Y \subseteq \bigcup_{\alpha\in F} U_\alpha$. (This is called the* Heine-Borel property*).*

The topological space X is said to be a compact space if X is compact as a subset of itself—namely, every open cover of X has a finite subcover.

Examples 2.44

(1) $(0,1)$*, the open unit interval, is* not *compact in $\mathbb{R}$ under the euclidean metric, even though it does admit* some finite *open covers, such as $\{(0,\frac{1}{2}),(\frac{1}{4},1)\}$. Instead if we let $U_n = (\frac{1}{n+2},\frac{1}{n})$, then $\{U_n : n \in \mathbb{N}\}$ is an open cover without any finite subcover.*

(2) *Let X be a metric space which is not complete. Then it is not compact. To see this, let $(x_n)_n$ be a Cauchy sequence which does not converge in X. Let $\overline{X}$ be a larger metric space containing X and in which the sequence converges to some point x_∞. Then for each $n \in \mathbb{N}$, let U_n be the set of points* in X, *which have distance greater than $\frac{1}{n}$ from x_∞.*

Lemma 2.45 *Let X be a compact space and $f : X \to \mathbb{R}$ be continuous. Then f is bounded and attains its maximum and minimum.*

There is much more that can be said for metric spaces.

Definition 2.46 (Totally bounded) *The subset Y of a metric space (X,d) is* totally bounded *if for every $\varepsilon > 0$ we can cover Y with finitely many balls of the form $B_\varepsilon(x)$.*

Every totally bounded set is bounded, but the converse is not true (for instance, an infinite set with the discrete metric is bounded but not totally bounded).

Theorem 2.47 (Compact metric spaces) *The following are equivalent for a subset Y of a metric space X:*

(1) *Y is compact.*

(2) *Y is complete and totally bounded.*

(3) *(Sequential compactness) Every sequence in Y has a subsequence that converges to a point of Y.*

(4) *Every continuous function from Y to $\mathbb{R}$ is bounded.*

(5) *Every continuous function from Y to $\mathbb{R}$ is bounded and attains its maximum.*

Sequential compactness is also called the *Bolzano-Weierstraß Property*. In particular, every compact subset of a metric space is closed and bounded. The converse is sometimes true:

Proposition 2.48 *(Heine-Borel) A subset of $\mathbb{R}^n$ is compact if and only if it is closed and bounded.*

An interesting consequence of Theorem 2.47 is that a complete, non-compact metric space is unbounded. It should be noted that "unbounded" only refers to the distance function—we will encounter complete, non-compact metric space that have nevertheless *finite volume*.

The following lemma shows that compactness is a topological property:

Lemma 2.49 *The continuous image of a compact space is compact.*

Proof If X is compact and $f : X \to Y$ continuous, let $C = f(X)$. Then for any open covering $\{U_\alpha : \alpha \in A\}$ of C, we can let $V_\alpha = f^{-1}(U_\alpha)$ for each $\alpha \in A$. From the definition of f being continuous, each V_α is open. Since $\{U_\alpha\}$ is a cover of $C = f(X)$, the $\{V_\alpha\}$ is an open cover of X. Since X is compact, there is some finite $F \subset A$ with $X \subset \bigcup_{\alpha\in F} V_\alpha$, which yields $C \subset \bigcup_{\alpha\in F} U_\alpha$. ∎

Lemma 2.50 *A closed subset of a compact space is compact.*

I'll conclude this section with an extremely useful fact:

Theorem 2.51 *If $f\colon (X,\mathscr{O}) \to (X',\mathscr{O}')$ is a continuous bijection of a compact space onto a Hausdorff space, then f is a homeomorphism.*

2.16 Separation axioms*

Metric spaces are a prime example of topological spaces. However, it turns out that the metric topology has many nice properties that topological spaces in general don't enjoy, such as the Hausdorff property. The following *separation properties* indicate different degrees of how well the open sets distinguish subsets of a topological space (beware that there are different versions of these in the literature):

T_0 If $x \neq y$, there exists an open U such that $x \in U$ but $y \notin U$,
or there exists an open U such that $x \notin U$ but $y \in U$.

T_1 If $x \neq y$, there exists an open U such that $x \in U$ but $y \notin U$.

T_2 (Hausdorff) If $x \neq y$, there exist open U and V such that $x \in U$, $y \in V$ and $U \cap V = \emptyset$.

T_3 (regular) $(X, \mathscr{O})$ is T_1 and for each closed $E \subseteq X$ and for each $x \in X \setminus E$
there exist open U and V such that $x \in U$, $E \subseteq V$ and $U \cap V = \emptyset$.

T_4 (normal) $(X, \mathscr{O})$ is T_1 and for all closed $E, F \subseteq X$ with $E \cap F = \emptyset$
there exist open U and V such that $E \subseteq U$, $F \subseteq V$ and $U \cap V = \emptyset$.

The above descriptions are best replaced by some pictures! Each separation property is a topological property. It is not difficult to show that $(X, \mathscr{O})$ satisfies T_1 if and only if $\{x\}$ is closed for each $x \in X$. Using this, one gets $T_4 \Rightarrow T_3 \Rightarrow T_2 \Rightarrow T_1 \Rightarrow T_0$.

Exercise 2.52 *The topology of any metric space satisfies T_4.*

Definition 2.53 *A topological space is* metrizable *if the topology is induced by some metric.*

Every metrizable space is normal by the above exercise, but not every normal space is metrizable. Necessary and sufficient conditions for metrizability are given by the Nagata–Smirnov metrization theorem, though there are also other sufficient conditions, for instance, the Urysohn metrization theorem. I won't give the statements since they involve the notion of a basis for a topology, which I haven't talked about.

Remark 2.54 *A* separable *space is a space that has a countable dense set. This has nothing to do with the separation axioms, and is rather unfortunate terminology.*

2.17 Manifold

(Related to [1] §12.5.)

Definition 2.55 (Manifold) *An n–*dimensional manifold *is a metric space with the property that every point has a neighbourhood that is homeomorphic to $\mathbb{R}^n$.*

The above definition is not the standard definition of a manifold. Usually one does not encounter a "metric space" in the definition, but a "paracompact Hausdorff space" or a "second countable Hausdorff space." All of these notions are equivalent if the space has at most a countable number of connected components, and there is no need to go into the relevant point-set topology. We will mainly be interested in the question of whether the quotient of a manifold under the action of a group of isometries is also a manifold.

Remark 2.56 *Thurston gives a more general definition of manifolds in [3] §3.1, and defines different structures on manifolds. Thurston's definition includes non-Hausdorff spaces that are locally homeomorphic to $\mathbb{R}^n$. Such spaces often arise naturally as quotient spaces of certain group actions or flows on a (Hausdorff) manifold.*

Using methods from algebraic topology, one can show that $\mathbb{R}^n$ is homeomorphic to $\mathbb{R}^m$ if and only if $n = m$. It follows from this *invariance of domain* that the dimension of a manifold is indeed well-defined. To save space, "n–dimensional manifold" is abbreviated to "n–manifold." It turns out that any n–manifold can be embedded in $\mathbb{R}^m$ for large enough m (depending on the manifold). For example, an open interval and the circle are 1–manifolds. The former can be embedded in $\mathbb{R}^1$ and the latter in $\mathbb{R}^2$. A class of examples generalising the standard embedding of the circle in the plane is given in §2.18.

Example 2.57 (1–manifolds) *Up to homeomorphism, the line and the circle are the only connected 1–manifolds. The circle is therefore the unique compact, connected 1–manifold.*

2.18 Regular level sets*

The following is a useful tool to obtain many examples of manifolds:

Proposition 2.58 (Regular level set is manifold) *Let $f\colon \mathbb{R}^{n+1} \to \mathbb{R}$ be a C^1 function and suppose $a \in \mathbb{R}$ is a regular value. Then the level set $f^{-1}(a)$ is an n–dimensional manifold.*

Proof Since $f^{-1}(a)$ is a (possibly not connected) subset of $\mathbb{R}^{n+1}$, we can turn it into a metric space by putting the induced metric on it. It remains to show that every point has a neighbourhood that is homeomorphic to $\mathbb{R}^n$. This follows from the Implicit Function Theorem. ■

Examples 2.59 (Two n–manifolds)

(1) *The n–sphere $\mathbb{S}^n$ is the level set $x_0^2 + \dots x_n^2 = 1$.*

(2) *The hyperboloid H^n is the level set $-x_0^2 + x_1^2 + \dots x_n^2 = -1$.*

Example 2.60 *$SL_n(\mathbb{R})$ is an (n^2-1)–dimensional manifold. This can be seen by viewing the set of all $n \times n$ matrices as $\mathbb{R}^{n^2}$. The determinant is a polynomial map, hence C^∞, with 1 as a regular value. So $\det^{-1}(1) = SL_n(\mathbb{R})$ is a manifold of the claimed dimension. Moreover, matrix multiplication $(M,N) \mapsto MN$ is a continuous map with respect to the product metric, and the inverse $M \to M^{-1}$ is also a continuous map. This is an example of a Lie group.*

2.19 Manifold-with-boundary

The following standard terminology (that may be confusing to start with) makes sense once one starts to consider manifolds *with boundary*.

Definition 2.61 *A compact manifold is called a* closed *manifold.*

Definition 2.62 (Manifold-with-boundary) *An* n–manifold-with-boundary *is a metric space with the property that every point has a neighbourhood that is homeomorphic to either $\mathbb{R}^n$ or the half-space*

$$\mathbb{R}^n_{\geq} = \{x \in \mathbb{R}^n \mid x_n \geq 0\}.$$

The circle is a closed 1–manifold, whilst the interval $[0,1]$ is a compact 1–manifold-with-boundary, and $(0,1]$ is a non-compact 1–manifold-with-boundary. In terms of their topology, each of these spaces is both open and closed, since it is regarded as a metric space in its own right.

The set of points whose neighbourhood is homeomorphic with the half-space is termed the *boundary* of the manifold-with-boundary, and the symbol ∂ is (again confusingly) used for the boundary. Whence $\partial[0,1] = \{0,1\}$ and $\partial(0,1] = \{1\}$, when both are regarded as manifolds-with-boundary.

If the boundary is empty, a manifold-with-boundary is simply a manifold as defined earlier. The hyphens will usually be dropped to allow statements such as "Let M be a manifold with (possibly empty) boundary…"

2.20 Surfaces

(Related to [1] §12.5.)

Definition 2.63 *A* surface (with boundary) *is a 2–manifold (with boundary).*

The following result was discovered by Jordan and Möbius in the 1860s, with missing details filled in by Dehn (1907) and Heegaard (1907). For pictures, terminology and a proof of the following theorem, see the account by George Francis and Jeff Weeks of Conway's Zero Irrelevancy Proof [20].

Theorem 2.64 (Classification of closed surfaces) *Every closed surface is homeomorphic to a sphere with a finite number of handles or a sphere with a finite number of cross caps.*

A sphere with handles is an orientable surface, whilst a sphere with cross caps is non-orientable. The next distinguishing feature is the *Euler characteristic*. There are different ways of defining it; the simplest definition is as follows. Every closed surface can be obtained by taking finitely many triangles and identifying their sides in pairs. Let F be the total number of triangles, E be the number of edges after identification, and V be the number of vertices after identification. Then

$$\chi(S) = V - E + F$$

is the Euler characteristic of S. It turns out that two closed surfaces are homeomorphic if and only if they have the same orientability class and the same Euler characteristic. Whence orientability class and Euler characteristic together are a *complete invariant* of a surface, i.e. an invariant that completely determines the topological type of a surface.

It is not surprising that every compact surface with boundary is obtained from a closed one by deleting a finite number of open discs. Whence, if one knows the number of boundary components, the orientability class and the Euler characteristic, one knows the topological type of a given surface with boundary.

Non-compact surfaces seem to be another story. For instance, the complement of a Cantor set in the plane is a surface. Amazingly, non-compact surfaces are also completely classified. Similar to the above classification of closed surfaces, Richards (1963) showed that *every* surface is obtained from a sphere by removing a (possibly infinite) number of points and open discs and then making suitable identifications between the boundaries of the discs to obtain handles or cross caps.

A complete invariant of such a surface is given by a nested triple of totally disconnected, compact, separable topological spaces. This can, in fact, be viewed as a nested triple of subsets of the standard Cantor set in $\mathbb{R}$. The triple of spaces roughly describes the orientablility class and genus (possibly infinite) plus what the surface looks like *at infinity*. This was done by Kerékjártó (1923) and Richards (1963). Every nested triple of totally disconnected, compact, separable topological spaces appears as the invariant of a non-compact surface.

Instead of computing and comparing the triple of spaces, Martin Goldman (1971) showed that to check homeomorphism of two non-compact surfaces, it suffices to compute some invariants from algebraic topology (the data for each surface consists of a vector space, a ring and three cup products).

In light of this, it is rather amazing that the *uniformisation theorem* (which is a generalisation of the Riemann mapping theorem and which was conjectured by Klein and Poincaré in the 1880s) implies that *every* surface has a Euclidean, spherical or hyperbolic structure. This was proved independently by Poincaré and Koebe in 1907. Bonahon states this geometrisation result as Theorem 12.15.

2.21 Three fundamental problems*

Homeomorphism Problem: Given two n–manifolds, determine whether they are homeomorphic.

As described above the Homeomorphism Problem has been solved completely in dimension 2.

An outline for the solution of the Homeomorphism Problem for a certain class of compact 3–manifolds with boundary (including all with non-empty boundary) was given by Haken in the 1960s, and the complete details were provided by Waldhausen (1968), Hemion (1976) and Matveev (2007). The problem for *closed* 3–manifolds has recently been completed with Perelman's work (2003), but it remains wide open for non-compact 3–manifolds.

In contrast, the Homeomorphism Problem cannot be solved in dimensions ≥ 4, even if one restricts to the class of all closed n–manifolds. This follows from a result in group theory: the problem of deciding whether two group presentations present isomorphic groups is undecidable. In these higher dimensions, one therefore tackles more restrictive sub-classes, where one may have a chance of succeeding.

Classification Problem: Determine a complete list of all n–manifolds (without duplications) together with an algorithm, which decides for any given n–manifold where it is on the list.

This problem has been solved for compact surfaces with boundary (and in particular for closed surfaces). Much research is devoted to solving this problem for *closed* 3–manifolds, but not even a good organising principle is known (though there are various competing approaches). Again, it cannot be solved in dimensions ≥ 4, and hence one tackles more restrictive sub-classes in these higher dimensions.

Algorithmic Classification Problem: Determine an algorithm which enumerates all n–manifolds (without duplications).

As before, this problem has been solved for compact surfaces with boundary, and it cannot be solved in dimensions ≥ 4. However, it is also solved for compact 3–manifolds (an outline is given in §2.22).

There is another reason that makes dimensions 2 and 3 special. Thurston [3] §3 talks about different structures that manifolds can have. The main structures in addition to the *topology (TOP)* induced by the metric are *differentiable (DIFF)* and *piecewise linear (PL)* structures. (The terms *differentiable manifold* and *smooth manifold* are used interchangeably, since any C^k structure ($k \geq 1$) has an (essentially unique) smoothing.) In dimensions 2 and 3 these distinctions break down: every topological 3–manifold can also be given an (essentially unique) differentiable and an (essentially unique) piecewise linear structure. This is not true in dimensions $n \geq 4$: There may be no differentiable or piecewise linear structures on a given n–manifold or there may be many inequivalent such structures. A famous example are the *exotic spheres* found by Milnor: There are inequivalent differentiable structures on $\mathbb{S}^7$, which all induce the standard topology on $\mathbb{S}^7$ (so one can go from one such structure to another by a homeomorphism, but not by a differentiable homeomorphism).

2.22 3–manifolds*

(Related to [1] §12.5.)

To describe the solution of the homeomorphism problem for closed 3–manifolds, one needs the following class of 3–manifolds. A *Seifert fibred space* is a compact 3–manifold (possibly with boundary) with a "nice" decomposition as a union of pairwise disjoint circles. The left picture in Figure 1 shows a solid cylinder foliated by intervals. Now identify the top disc with the bottom disc, such that the interval in the centre has its endpoints identified. This still allows freedom to rotate the disc about its centre before identification. We require that every interval in the resulting solid torus is part of a circle, or, equivalently, that the angle of rotation is a rational multiple of π. If no rotation is applied, then all circles travel in parallel, but rotation by 2π gives a different result since the circle in the centre is now distinguished, having all other circles spiral around it. The picture in the middle of Figure 1 shows the result after rotation and identification. The requirement for a Seifert fibered space is that *every* circle has a neighbourhood that is obtained in this manner from a solid cylinder.

A natural equivalence relation on a Seifert fibred space is defined by declaring points equivalent if and only if they lie on the same circle. The set of equivalence classes is a surface, but it has some extra structure that keeps track of how circles spiral around each other. Compactness implies that there are only finitely many special circles which don't travel parallel with all of their neighbours. The surface with this extra data is an *orbifold*, and the Seifert fibred space is completely encoded by the *Seifert invariants*, which record the topology of the underlying surface and the information about the special points. Seifert (1933) defined these special 3–dimensional spaces and solved the classification and homeomorphism problems for them.

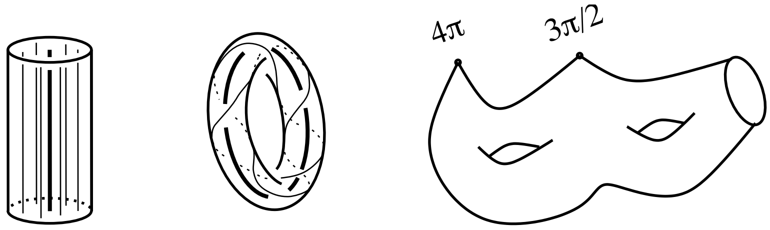

Figure 1: A Seifert fibred space is an orbifold's worth of circles

Theorem 2.65 (Almost the Geometrisation Theorem) *Every closed, orientable 3–manifold has a* unique decomposition *along spheres and tori, such that each piece is either a* Seifert fibred space *or a* hyperbolic 3–manifold of finite volume.

The above theorem combines the work of Kneser (1927), Milnor (1962), Jaco & Shalen (1979), Johannson (1979), Thurston (1979), Hamilton (1981) and Perelman (2003). I call it *Almost the Geometrisation Theorem* since it does not give the standard decomposition of the manifold into geometric pieces, but the latter can be derived from the decomposition in the theorem. The proof of the theorem has two parts; one is topological and the other geometric. The *topological part* consists of the decomposition first along spheres (Kneser and Milnor), and then along tori and Klein bottles (Jaco & Shalen and independently Johannson). Thurston (1982) and Perelman (2006) won the Fields medal for the mathematics behind the *geometric part* of the proof. The *Almost the Geometrisation Theorem* implies the *Poincaré conjecture*, which states that every simply connected, closed 3-manifold is homeomorphic to the 3-dimensional sphere. Intuitively, *simply connected* means that every closed loop in the space can continuously be shrunk to a point. See §4 *"Some algebraic topology"* for a precise definition.

Almost concurrently to the above work, much effort has been put into devising algorithms to find the decomposition as well as the structures on the pieces. The foundation is a result by Moise which states that every compact 3–manifold with boundary can be decomposed into a finite number of tetrahedra (just as a surface can be decomposed into triangles). The following result summarises work of Moise, Bing, Haken, Jaco, Oertel, Tollefson, Rubinstein, Stocking, Thompson, Li, Casson, Manning and Sela.

Corollary 2.66 *Given a closed, orientable 3–manifold, there are algorithms to find its topological decomposition, and to construct the structures on the pieces.*

The solution to the homeomorphism problem follows from the decomposition theorem: Given two 3–manifolds, M and N, compute the topological decomposition. If the number of pieces is different, M and N are not homeomorphic. Otherwise construct the structures on the pieces, and check whether they match up in pairs. For the Seifert fibered pieces, this is easily done using the above mentioned Seifert invariants. For hyperbolic pieces, this is a little more involved, but there are many different techniques that can be used. If the pieces don't match up in pairs, M and N cannot be homeomorphic. Otherwise check whether they were glued up in the same way to give the manifolds M and N. If they were not, then M and N are not homeomorphic, and otherwise they are.

Given the fact that every compact 3–manifold can be decomposed into a finite number of tetrahedra, a *cheapological trick* (as it was coined by Haken) now gives a solution to the algorithmic classification problem. The trick consists of slowly building a census of 3–manifolds that can be made up of a single tetrahedron, then two tetrahedra, then three, and so forth. Every time a new manifold has been created, one checks whether it is homeomorphic to any of the ones that are already in the list.

Amongst the above fundamental problems, the only one that remains open for closed 3–manifolds is the classification problem. Since Seifert fibred spaces are completely classified, for a complete solution it remains to classify the hyperbolic 3–manifolds of finite volume. At the time of writing, this problem is wide open, and it is one of many reasons why the study of hyperbolic 3–manifolds is an extremely active area of research.

Bibliography

The literature listed under *Bed-time reading* and *Sources* has been used in preparing the lectures or supplementary notes, or was intended to be used. References to additional research papers and survey articles will be given during the course.

Textbook

[1] Francis Bonahon: *Low-Dimensional Geometry: From Euclidean Surfaces to Hyperbolic Knots*. Student Mathematical Library, 49. IAS/Park City Mathematical Subseries. American Mathematical Society, Providence, RI; Institute for Advanced Study (IAS), Princeton, NJ, 2009.

Other main references

[2] Walter D. Neumann: *Notes on geometry and 3-manifolds*. With appendices by Paul Norbury. Bolyai Soc. Math. Stud., 8, Low dimensional topology (Eger, 1996/Budapest, 1998), 191267, János Bolyai Math. Soc., Budapest, 1999. Available from: www.math.columbia.edu/∼neumann/preprints/

[3] William P. Thurston: *Three-Dimensional Geometry and Topology*. Princeton Mathematical Series, 35. Princeton University Press, Princeton, NJ, 1997.

[4] Stephan Tillmann: *"Geometry & Groups"* (supplementary notes)

Bed-time reading

[5] David Hilbert and S. Cohn-Vossen: *Geometry and the Imagination*. Chelsea Publishing, New York, 1952.

[6] David Mumford, Caroline Series and David Wright: *Indra's pearls. The vision of Felix Klein.* Cambridge University Press, New York, 2002.

[7] Jeffrey Weeks: *The shape of space*, Monographs and Textbooks in Pure and Applied Mathematics, 249. Marcel Dekker, New York, 2002.

Geometry games and software

[8] Marc Culler and Nathan Dunfield: *SnapPy*, freely available at http://www.math.uic.edu/t3m/SnapPy/doc/

[9] Daniel J. Heath: *Geometry playground*, a ruler and compass Java application for several 2–dimensional geometries, freely available at http://www.plu.edu/∼heathdj/java/

[10] Masaaki Wada: *OPTi*, a Macintosh application which visualizes quasi-conformal deformations of the once-punctured-torus groups, freely available from http://delta.math.sci.osaka-u.ac.jp/OPTi/index.html

[11] Jeffrey Weeks: *Topology and Geometry Software*, freely available at www.geometrygames.org

Sources

[12] Roger C. Alperin: *An elementary account of Selberg's lemma*, L'Enseignement Mathématique 33 (1987) 269–273.

[13] Alan F. Beardon: *The geometry of discrete groups.* Springer-Verlag, New York, 1995.

[14] Riccardo Benedetti and Carlo Petronio: *Lectures on hyperbolic geometry*. Springer-Verlag, Berlin, 1992.

[15] Michel Boileau, Sylvain Maillot and Joan Porti: *Three-dimensional orbifolds and their geometric structures.* Panoramas et Synthèses, 15. Société Mathématique de France, Paris, 2003.

[16] Martin R. Bridson and André Haefliger: *Metric spaces of non-positive curvature.* Springer-Verlag, Berlin, 1999.

[17] Herbert Busemann and Paul J. Kelly: *Projective geometry and projective metrics.* Academic Press, New York, 1953.

[18] Daryl Cooper, Craig D. Hodgson and Steven Kerckhoff: *Three-dimensional orbifolds and cone-manifolds.* MSJ Memoirs, 5. Mathematical Society of Japan, Tokyo, 2000.

[19] Gerald B. Folland: *Real analysis. Modern techniques and their applications.* Second edition. John Wiley & Sons, New York, 1999.

[20] George Francis and Jeffrey Weeks: *Conway's ZIP Proof.* Amer. Math. Monthly 106 (1999), no. 5, 393399.

[21] William M. Goldman: *Convex real projective structures on compact surfaces*. Journal of Differential Geometry 31 (1990) 791–845.

[22] Michael Kapovich: *Hyperbolic manifolds and discrete groups.* Birkhäuser, Boston, MA, 2009.

[23] Svetlana Katok: *Fuchsian groups*. Chicago Lectures in Mathematics. University of Chicago Press, Chicago, 1992.

[24] Marc Lackenby: *Topology and Groups*. Lecture notes. Available from: http://people.maths.ox.ac.uk/lackenby/

[25] Albert Marden: *Outer circles. An introduction to hyperbolic 3-manifolds.* Cambridge University Press, Cambridge, 2007.

[26] Bernard Maskit: *On Poincaré's Theorem for Fundamental Polygons*. Advances in Mathematics 7 (1971) 219–230.

[27] John Milnor: *Hyperbolic geometry: The first 150 years*. Bulletin (New Series) of the American Mathematical Society 6 (1982) 9–24.

[28] James Munkres: *Topology.* Second edition. Prentice-Hall, Upper Saddle River, 2000.

[29] Tristan Needham: *Visual Complex Analysis*. The Clarendon Press, Oxford University Press, New York, 1997.

[30] Wayne Patty: *Foundations of topology.* Second edition. Jones and Bartlett Publishers, Boston, MA, 2009.

[31] John G. Ratcliffe: *Foundations of hyperbolic manifolds*. Springer, New York, 2006.

[32] Richard Evan Schwartz: *Mostly surfaces*. Student Mathematical Library, 60. American Mathematical Society, Providence, RI, 2011.

[33] Peter Scott: *The geometries of 3-manifolds*. Bulletin of the London Mathematical Society 15 (1983) 401–487.

[34] Herbert Seifert and William Threlfall: *Lehrbuch der Topologie*. AMS Chelsea Publishing, 1934 (reprinted 2003).

[35] John R. Stallings: *Topology of finite graphs.* Invent. Math. 71 (1983), no. 3, 551565.

[36] William P. Thurston: *The Geometry and Topology of Three-Manifolds*. Lecture notes, Princeton University 1980. Available at http://library.msri.org/books/gt3m/ (Electronic version 1.1 March 2002)