Lecture 2: Basics

Topics in Computational Topology: An Algorithmic View

Today we will introduce some basic concepts in topology. This lecture and next two lectures we will still stay in continuous domain. After that, we will change gear to discrete domains, handling discrete objects (especially simplicial complexes).

1 Topological Spaces

In a very abstract manner, a *topological space* is a set X endowed with a *topological structure* (a *topology*) $\mathcal{T}$ such that the following conditions are satisfied:

1. Both the empty set and X are elements of $\mathcal{T}$.
2. Any union of arbitrarily many elements of $\mathcal{T}$ is an element of $\mathcal{T}$.
3. Any intersection of finitely many elements of $\mathcal{T}$ is an element of $\mathcal{T}$.

In other words, it is a set equipped with set of subsets. In particular, each set A in $\mathcal{T}$ is said to be *open*. Its complement $\mathsf{X} \setminus A$ is said to be *closed*. (Note that a set can be both open and closed.) That is, a set A is *closed* if its complement is open. The *closure* of a set A is the minimal closed set from X containing A. Given the same space (set) X, one can defined different set system $\mathcal{T}$, which will then in turn lead to different topology.

Now this looks a bit too alien and very abstract. Let us look some examples.

1. Let X be a set of n elements $\mathsf{X} = \{x_1, \ldots, x_n\}$, and let $\mathcal{T} = 2^{\mathsf{X}}$. $< \mathsf{X}, \mathcal{T} >$ forms a topology. This is a *discrete* topology [1]. It is a simple topology, but somewhat meaningless.
2. Given a graph $G = (V, E)$, let $\mathsf{X} = \{V, E\}$ and define an open set to be any subgraph of G. This topology is also a discrete topology.
3. Let us now try to rephrase everything in the metric space. First, let us consider the Euclidean space $\mathbb{R}^d$. A set $A \subseteq \mathbb{R}^d$ is *open* if for any point $x \in X$, there exists some $\varepsilon > 0$ such that any y with $\|x - y\| < \varepsilon$ is also in A. Intuitively, A is open if at any point inside, we can move in arbitrary direction while still staying inside A. (Consider the more familiar *open intervals* on a line.) The requirements in previous definitions of topology are easily satisfied. Hence we have a topological space of $\mathbb{R}^d$ equipped with this open set. This definition of open set can be extended to any metric space[2] by replacing the distance $\|x - y\|$ with whatever metric distance $d(x, y)$ we need, which will give us a natural topology for any metric space. It is easy to check that the following set of open balls form the basis of the standard topology of a metric space described above. $\mathcal{W} := \{B_\varepsilon(x) \mid x \in \mathsf{X}\}$ where $B_\varepsilon(x) = \{y \in \mathsf{X} \mid d(x, y) < \varepsilon\}$. The natural topology defined on a metric space perhaps is perhaps the most important and most common topological space.

[1] A topology $< \mathsf{X}, \mathcal{T} >$ is *discrete* if $\mathcal{T} = 2^{\mathsf{X}}$.

[2] Recall a metric space is a space equipped with a distance $d(x, y)$ defined for any two elements $x, y \in \mathsf{X}$ such that the following conditions are satisfied: (1) $d(x, y) \geq 0$, (2) $d(x, y) = 0$ iff $x = y$, (3) $d(x, y) = d(y, x)$, (4) $d(x, y) \leq d(x, z) + d(z, y)$.

Note that for the real line, the intersection of the intervals $(-\frac{1}{k}, +\frac{1}{k})$, for all integers $k \geq 1$, is the point 0. This is not an open set. This illustrates the need for the restriction to finite intersections.

Finally, suppose that we have a topological space $< \mathsf{X}, \mathcal{T} >$. Given a subset $\mathsf{Y} \subseteq \mathsf{X}$, it has a natural topology on it which is inherited from $\mathcal{T}$, denoted by $\mathcal{T}_Y$ defined as follows: open sets in $\mathcal{T}_Y$ is the intersection between open sets in $\mathcal{T}$ and Y. This is called the *topology on* Y *induced from* $< \mathsf{X}, \mathcal{T} >$. For example, when we talk about topology for a surface S embedded in $\mathbb{R}^3$, we in fact mean the topology on S induced from $\mathbb{R}^3$. From now on, I will often omit the explicit reference of $\mathcal{T}$ and simply talk about a topological space X when the choice of $\mathcal{T}$ is clear. (In fact, we will mostly talk about the topology induced from a Euclidean space in this class.)

Remark. The topology (as well as the induced topology) in Euclidean space is the most common topological space one will encounter. The definition of topological space as sets of subsets may seem un-natural at first. In particular, the definitions of *open* and *closed* sets may be non-intuitive. Recall that we have said earlier that topology is about *connectivity* and about how the input space is put together from its subsets. Intuitively, this is captured by the open set and their union and intersections. Now when we compare two topological spaces, we need a map between them, and we need a language to say that the two spaces are connected in the same way using this map. The language for this purpose is *continuity*, which is one of the most important concepts in not just topology, but mathematics.

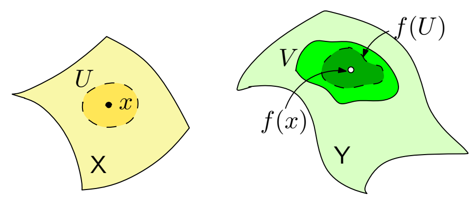

Continuity. A *neighborhood* of a point $x \in \mathsf{X}$ is simply a subset of X that contains some open set U such that $x \in U$. A function $f : \mathsf{X} \to \mathsf{Y}$ is *continuous at* $x \in \mathsf{X}$ if for any neighborhood V of $f(x)$, there is a neighborhood U of x such that $f(U) \subseteq V$. See the right figure for an illustration. Function f is *continuous* if it is continuous at all points in X. Note that for the space $\mathsf{X} = \mathbb{R}$ endowed with the standard topology and for a function $f : \mathbb{R} \to \mathbb{R}$, this definition coincides with the standard (ε, δ) definition from Calculus. (Consider the example where the function f has a jump at x.) An alternative, perhaps simpler definition for continuity is that: a function $f : \mathsf{X} \to \mathsf{Y}$ is *continuous* if any open set in Y has an open pre-image. For discrete topology, we have the following result:

Claim 1.1 *Any map $f : \mathsf{X} \to \mathsf{Y}$ of a space X with discrete topology is continuous.*

Homeomorphism. The most important concept to study topology is homeomorphism. Given two topological spaces X and Y, there is a *homeomorphism* between them, denoted by $\mathsf{X} \approx \mathsf{Y}$, if there is continuous function $f : \mathsf{X} \to \mathsf{Y}$ such that f has an inverse map $f^{-1} : \mathsf{Y} \to \mathsf{X}$ such that $f \circ f^{-1} = id$, which is also continuous. The requirement that the inverse is also continuous is important. For example, look at the example $f : [0, 2\pi) \to \mathbb{S}_1$ with $f(x) = (\sin x, \cos x)$. f does not induce a homeomorphism between an interval and a circle as its inverse it not continuous.

Informally, that the inverse exists means *bijection* which roughly means that we have a one-to-one correspondence between the two sets. That both f and its inverse are continuous mean that there is a bijection between open sets from X and open sets from Y, and they are all connected in the same way.

Now let us look at a few examples.

(1) an open disk and $\mathbb{R}^2$. The explicit mapping is $f(x) = \frac{x}{1 - \|x\|}$. (In fact, this map establishes a homeomorphism between the open d-ball and $\mathbb{R}^d$ for any $d > 0$.)

(2) Sphere and a tetrahedron. (From the center shoot a ray in all directions, and it intersects both tetrahedron and the sphere. That is the mapping f. (In fact, this map works for any two convex body of the same dimension.)

(3) Sphere with the north pole point removed and $\mathbb{R}^2$. (Again, shoot rays from north pole.)

We do not always need to find an explicit mapping to see that two spaces are homeomorphic. Intuitively, if one can deform from either one to the other without breaking and inserting, then they are homeomorphic. (Recall the examples we had from Lecture 1.)

2 Connectedness

A *path* is a continuous function from the unit interval, $\gamma : [0, 1] \to \mathsf{X}$. A path is *closed* if $\gamma(0) = \gamma(1)$. A closed path is also called a *closed curve*. (An alternative way to define a close curve is a continuous function from the unit circle $\gamma : \mathbb{S}_1 \to \mathsf{X}$ (where $\mathbb{S}_1 = \{x \in \mathbb{R}^2 \mid \|x\| = 1\}$).) A path γ is *simple* if it is injective (other than potentially for $\gamma(0) = \gamma(1)$). A simple path does not have self-intersections. Often, we abuse the notation slightly and use a path to both refer to the map γ itself as well as its image $\gamma([0, 1]) \subset \mathsf{X}$ in X.

Definition 2.1 *A topological space* X *is* path connected *if for every* $x, y \in \mathsf{X}$, *there is a path connecting them. A topological space* X *is* connected *if there does not exist two disjoint non-empty open sets* $X_1, X_2 \subset \mathsf{X}$ *such that* $X_1 \cup X_2 = \mathsf{X}$. *A maximal connected subset of* X *is called a* connected component *of* X.

The definition of connectedness is slightly weaker than path-connectedness. However, for most spaces we will ever encounter, they are the same.

Theorem 2.2 (Jordan Curve Theorem) *Removing the image of a simple closed curve from* $\mathbb{R}^2$ *leaves two connected components. The unbounded one is called the* outside *one, and the other one bounded* inside.

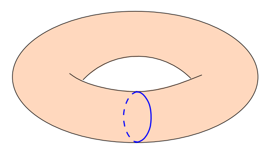

The Jordan Curve Theorem can be extended to the 2-sphere $\mathbb{S}_2$. However, it does not hold for any other surface. Consider an example on the torus – a closed curve around the handle of the torus does not disconnect the surface at all! (See the right figure for an illustration.) Later in this class we will see exactly what types of curves will have this "non-separating" property.

3 Manifolds

First, some notations. $\mathbb{B}^o_d = \{x \in \mathbb{R}^d \mid \|x\| < 1\}$ is the open d-ball. We use $\mathbb{B}_d = \{x \in \mathbb{R}^d \mid \|x\| \leq 1\}$ to denote the closed d-ball. The d-sphere is the boundary of a $d+1$ ball. That is, $\mathbb{S}_d = \{x \in \mathbb{R}^{d+1} \mid \|x\| = 1\}$. What is $\mathbb{S}_0$, $\mathbb{S}_1$, etc? Recall that we have shown earlier that $\mathbb{B}^o_d \approx \mathbb{R}^d$; i.e, $\mathbb{B}^o_d$ is homeomorphic to $\mathbb{R}^d$. If we remove a single point from $\mathbb{S}_d$, what do we get? (We obtain a space homeomorphic to $\mathbb{B}^o_d$ and $\mathbb{R}^d$.)

A *d-manifold (without boundary)* is a topological space M such that each point $x \in \mathsf{M}$ has a neighborhood homeomorphic to $\mathbb{R}^d$ (i.e, to $\mathbb{B}^o_d$). A *d-manifold with boundary* is a topological space M such that each point $x \in \mathsf{M}$ either has a neighborhood homeomorphic to $\mathbb{R}^d$, or has a neighborhood homeomorphic to half-space $\mathbb{R}^d_+$. The collection of latter points are called the *boundary* of the manifold.

Theorem 3.1 (Boundary of A Manifold) *The boundary of a* m*-manifold is a* $(m-1)$*-dimensional manifold (possibly with multiple connected components).*

Basically, manifolds are those that locally it looks like the Euclidean space. Any closed curve is a 1-manifold. A curve is 1-manifold with boundary. A surface is 2-manifold, like a sphere, and a torus. In graphics, a solid is 3-manifold with boundary. Now some examples of non-manifolds: (1) a curve with branches, (2) two balls with a pinch point in the middle, (3) how about a sphere with a hole ? (this is manifold with boundary!) (4) möbius strip (still manifold with boundary), (5) a book.

Embedding and immersion. Note that we define m-manifold in an abstract, local manner. A closed curve is a 1-manifold no matter which space it is in. Now we can imagine putting an input m-manifold M in certain Euclidean space $\mathbb{R}^d$. This is simply a mapping $\Phi : \mathsf{M} \rightarrow \mathbb{R}^d$.

Definition 3.2 *We say that Φ is an* immersion *if locally for every point $x \in \mathsf{M}$, there is a small neighborhood $U(x)$ such that Φ induces an homeomorphism from $U(x)$ to $\Phi(U(x))$. Φ is an* embedding *if it induces a homeomorphism between M and $\Phi(\mathsf{M})$.*

The dimension m is called the *intrinsic dimension* of M. The space that M is embedded in is called the *ambient space* and its dimension is *ambient dimension*. For example, in our real life, we usually consider 1- or 2-manifolds embedded in $\mathbb{R}^3$. Note that ambient dimension is at least the intrinsic dimension.

Compactness.

Definition 3.3 *A topological space X is* compact *if any cover of X by open sets (i.e, a collection of open sets whose union coincides with X) admits a finite sub-cover.*

Definition 3.4 *A subset X of a metric space Ω with endowed metric d is* bounded *if it is contained in a ball of finite radius, i.e. if there exists $u \in \Omega$ and $r > 0$ such that for all $x \in X$, we have $d(x, u) < r$. Ω is a bounded metric space (or d is a bounded metric) if Ω is bounded as a subset of itself.*

Theorem 3.5 (Compact Euclidean Domains) *A subset of the Euclidean space $\mathbb{R}^d$ is compact if and only if it is closed and bounded.*

Hence, the sphere $\mathbb{S}_d$ is compact, but the interval $(0, 1)$ is not. It is intuitive to see that homeomorphism preserves compactness.

Acknowledgment. We thank Dong Du for pointing out several imprecise statements in an earlier version of this lecture note.