APPLIED TOPOLOGY LECTURE NOTES

CHAD GIUSTI

UPDATED: SEPTEMBER 6, 2017

Introduction: Thinking Topologically

WHAT IS TOPOLOGY? Here, we will take the slightly unorthodox (but useful) view that *topology* is the mathematics of qualitative notions of similarity. Possibly the simplest example of such a notion is a fixed group of objects with a binary notion of pairwise similarity: for any pair of objects, the CEpair is either similar or dissimilar. This type of data is exceedingly common and is called a *graph*.

Graphs

We begin by fixing some notation that will get a lot of use.

Definition. Let S be a set.

- Write |S| for the *cardinality* of the set S. If S is a finite set, this is just the number of elements in |S|.
- The set of *ordered k-tuples of elements of S* is written $S^{\times k} = S \times S \times \cdots \times S$, and elements are written $(s_1, s_2, \ldots, s_k)$.
- The set of *unordered k-element subsets of S* is denoted $\binom{S}{k}$. Denote by $s_1 s_2 \ldots s_k$ (or any permutation thereof) the element $\{s_1, s_2, \ldots s_k\}$.
- The *power set of S* is the collection of all subsets of S, written 2^S.[1]

[1] This notation is fantastic because $2^{|S|} = \sum_{k=0}^{|S|} \binom{|S|}{k}$, and $2^S = \bigcup_{k=0}^{|S|} \binom{S}{k}$.

Definition. An *(abstract, simple) graph* Γ is a pair of sets $\Gamma = (V, E)$ with $E \subseteq \binom{V}{2}$.

Throughout these notes, unless explicitly stated otherwise, all graphs are simple and *finite*; that is, $|V| < \infty$.

In our example above, the vertices $v \in V$ are the objects under consideration, and the presence of an edge $vv' \in E$ indicates that v and v' are similar. Without any further information, this is a very general notion of similarity: knowing that $v_1 v_2 \in E$ and $v_2 v_3 \in E$ tells us nothing about whether v_1 is similar to v_3.

Examples.

- $V =$ people, $E =$ self-identified friends (social media)
- $V =$ locations, $E =$ connecting paths (map software)

- V = products, E = purchased together (online shopping)
- $V = \varnothing$, $E = \varnothing$ (the empty graph)
- $V = \{p\}$, $E = \varnothing$ [2]
- a) $V_1 = \{p, q\}$, $E_1 = \varnothing$
  b) $V_2 = \{p, q\}$, $E_2 = \{pq\}$ [3]
- a) $V_1 = \{p, q, r, s\}$, $E_1 = \{pq, ps, qr, rs\}$
  b) $V_2 = \{w, x, y, z\}$, $E_2 = \{wx, wz, wy, yz\}$
  c) $V_3 = \{i, j, k, \ell\}$, $E_3 = \{ij, ik, i\ell, k\ell\}$ [4]

[2] Is this the only graph with one vertex? How might it fail to be unique?

[3] These are distinct graphs on the same vertex set $V_1 = V_2$. How can we tell them apart? The easiest way is to compare number of edges: $|E_1| \neq |E_2|$. Is this a complete list of graphs on two vertices?

[4] Are these the same? $|V_1| = |V_2| = |V_3|$ and $|E_1| = |E_2| = |E_3|$, so counting elements won't work here. Looks like it's time for a more sophisticated approach.

How do we compare objects mathematically to determine "sameness"? The most fundamental and familiar situation is when our objects are sets, say S and T. If the sets are "the same", to each element in one of them, say S, we can assign a unique corresponding element in T, and in doing so exhaust the elements of T. Formally, doing this is performed by defining a function from S to T. In some settings functions come naturally from semantics: if S is a set of people and T is a set of dates, natural assignments include things like birthdays.

On the other hand, in the absence of clear semantics, we are often interested in only the *existence* of a function which gives us such a correspondence. Two of the most useful, which we will see repeatedly are *injectivity* (sometimes called *being one-to-one*) and *surjectivity* (or *being onto*). The existence of an injective function $f : S \to T$ implies that $|S| \leq |T|$, and the existence of a surjective function $g : S \to T$ implies that $|S| \geq |T|$. Finding such a pair of functions or a single *bijective* (both injective and surjective) function is sufficient to conclude that $|S| = |T|$. Since the only inherent structure in a set are the elements themselves, this is sufficient to tell us that the structure is (in an abstract sense) the same: S and T are *isomorphic*, written $S \cong T$. [5]

[5] We've already used the term cardinality when refering to sets without thinking too deeply about its meaning, but to be rigorous we need to define it in terms of *isomorphism classes* of sets.

It bears repeating that this notion of isomorphism is abstract: any data besides the fact that the sets contain elements is lost. In most cases, we will need to be cautious to ensure that other structures we build on top of our sets are maintained. In our running example, graphs consist of two sets, V and E, but the elements of the set E have meaning endowed by V. Thus, if we define a function between vertex sets, we also get an interaction between edges for "free".

Definition. Let $\Gamma_1 = (V_1, E_1)$ and $\Gamma_2 = (V_2, E_2)$ be graphs. A function $\varphi : V_1 \to V_2$ *induces* a relation[6] $\widetilde{\varphi} : E_1 \to E_2$ by $\widetilde{\varphi}(vv') = \varphi(v)\varphi(v')$.

[6] Recall that a *relation* R between sets S and T is just a collection of ordered pairs $R \subseteq S \times T$.

Note that this is not necessarily a function: if $vv' \in E_1$ but $\varphi(v)\varphi(v') \notin E_2$, the relation $\widetilde{\varphi}$ is undefined on vv'. However, if the relation *is* defined for every element of E_1, the induced relation is a function, and this pair is our notion of "structure preserving function" or *morphism* between graphs.

Definition. Let $\Gamma_1 = (V_1, E_1)$ and $\Gamma_2 = (V_2, E_2)$ be graphs, and $\varphi : V_1 \to V_2$ a function for which the induced relation $\widetilde{\varphi}$ is a function. The the pair of functions $(\varphi, \widetilde{\varphi})$ is a *(graph) homomorphism*, written $\Phi : \Gamma_1 \to \Gamma_2$.

So, when are two graphs "the same"? When the underlying vertices can be identified in such a way that the edge structure is also preserved.

Definition. Let $\Gamma_1 = (V_1, E_1)$ and $\Gamma_2 = (V_2, E_2)$ be graphs. If there exists a homomorphism $\Phi : \Gamma_1 \to \Gamma_2$ so that $\varphi : V_1 \to V_2$ and $\widetilde{\varphi} : E_1 \to E_2$ are both bijections, then Γ_1 is *(graph) isomorphic* to Γ_2, written $\Gamma_1 \cong \Gamma_2$.

Graph isomorphisms preserve all of the structure inherent in graphs, so it is often convenient and sensible to forget which specific graph we're discussing and instead move freely between isomorphic graphs. The collection of all graphs isomorphic to a fixed graph Γ is called its *isomorphism class*. We will often adopt the highly abusive convention of saying "the graph Γ" and meaning "the graph isomorphism class of Γ" when this won't create confusion.
Back to our last example:

a) $V_1 = \{p, q, r, s\}$, $E_1 = \{pq, ps, qr, rs\}$
b) $V_2 = \{w, x, y, z\}$, $E_2 = \{wx, wz, xy, yz\}$
c) $V_3 = \{i, j, k, \ell\}$, $E_3 = \{ij, ik, i\ell, k\ell\}$

There are $4! = 24$ possible bijections[7] $\varphi : V_1 \to V_2$. Some of these induce functions $\widetilde{\varphi} : E_1 \to E_2$ and others don't. If we choose a target for p, say $\varphi(p) = x$, then our choices are constrained. Observe that p is implicated in two edges, pq and ps. That tells us that $\varphi(p)\varphi(q) = x\varphi(q)$ and $\varphi(p)\varphi(s) = x\varphi(s)$ must be elements of E_2. Since the edges in E_2 containing x are wx and wz, we must have $\varphi(q) = x$ or z and $\varphi(s)$ as the other. Either choice forces $\varphi(r) = y$, and the resulting function induces a graph isomorphism.

[7] There are $n!$ possible bijections between two sets with n elements. As you can imagine, the problem of finding a graph isomorphism can become very difficult – indeed, it is in class NP.

However, there are no graph isomorphisms between the first (or second) and third graphs. Why? Whatever vertex from V_1 gets sent to vertex i must be implicated in three edges in order for the induced map on edges to be surjective, since the map must hit all three of ij, ik and $i\ell$. However, there are no such vertices in the first graph, so there is no choice of bijection of vertices which induces a graph isomorphism.

Definition. Let $\Gamma = (V, E)$ be a graph and $v \in V$. The *degree* of v is the number of edges $e \in E$ which contain v. The *degree sequence* of a Γ is the list of the degrees of all vertices of Γ, sorted in decreasing order.

In order, the degree sequences of our graphs above are $(2,2,2,2)$, $(2,2,2,2)$, and $(3,2,2,1)$.

Lemma 1. *Let Γ_1 and Γ_2 be graphs with degree sequences (d_i) and (d'_i) respectively. If $\Gamma_1 \cong \Gamma_2$, then $(d_i) = (d'_i)$.*

Proof. Exercise. □

The converse of this statement is that the degree seqeunce of a graph is an *isomorphism invariant*: if two graphs have different degree sequences, they cannot be isomorphic.

Example.
a) $V_1 = \{p,q,r,s,t,u\}$, $E_1 = \{pq, pr, qr, st, su, tu\}$
b) $V_2 = \{a,b,c,d,e,f\}$, $E_2 = \{ab, ac, bd, ce, df, ef\}$

Are these isomorphic? There are $6! = 720$ possible bijections between V_1 and V_2, so it's going to be tedious to enumerate them. Every vertex is degree two, so we can't rely on simple counting arguments to cheat. Let's consider a more interesting property.

Definition. The *line graph of order n* is the graph

$$L_n = (\{v_1, \ldots v_n\}, \{v_1v_2, v_2v_3, \ldots, v_{n-1}v_n\}.$$

Definition. Let $\Gamma = (V,E)$ be a graph with $a, b \in V$. A *path of length $n+1$ from a to b* is a homomorphism $P : L_n \to \Gamma$ so that $p(v_1) = a$ and $p(v_n) = b$. There is a *path from a to b in* Γ if there is a path of length $n+1$ for some n. A path P is *simple* if it is injective on vertices. A *cycle* is a path P which is injective on vertices except that $p(v_1) = p(v_n)$.

Observe that this definition coincides with the "regular" definition of a path in a graph being a sequence of vertices with consecutive vertices joined by edges. However, it makes use of a *reference object*, the line graph, which is simple to understand. Exploring the existence of maps out of or into reference objects are one of the fundamental ways we will deconstruct complicated objects we encounter.

Definition. There is a unique partition of V into subsets $\mathcal{C} = \{C_1, C_2, \ldots, C_n\}$ so that for every i and pair of vertices $v, v' \in C_i$, there is a path from v to v′, but for any $v \in C_i$ and $v' \in C_j$, $i \neq j$, there is no path from v to v′. The sets $\mathcal{C}$ are called the *(path) components* of Γ, and a graph with exactly one path component is called *(path) connected*.

Definition. Let $\Gamma = (V,E)$ be a graph. Let $V' \subseteq V$ and $E' \subseteq E \bigcap \binom{V'}{2}$. The graph $\Gamma' = (V', E')$ is called a *subgraph* of Γ. If $E' = E \bigcap \binom{V'}{2}$, we say Γ' is the subgraph of Γ *induced* by V'.

Lemma 2. *Let* Γ_1 *and* Γ_2 *be graphs with path components* $\mathcal{C} = \{C_1, C_2, \ldots, C_n\}$ *and* $\mathcal{D} = \{D_1, D_2, \ldots, D_m\}$ *respectively. If* $\Gamma_1 \cong \Gamma_2$*, then* $n = m$*. Further, there is a bijection* $\psi : \mathcal{C} \to \mathcal{D}$ *so that, for each* $i = 1, \ldots, n$*, the subgraph of* Γ_1 *induced by* C_i *is isomorphic to the subgraph of* Γ_2 *induced by* $\psi(C_i)$*.*

Proof. Exercise. □

In light of this observation, it is common practice to restrict attention to the study of path connected graphs, as we can always decompose a graph into path-connected components, which essentially do not interact.

How many path components do the graphs in our example have? Observe that in E_1, there is no path from p to t: from p one can reach q or r, and each of these has an edge to the other, but there are no edges to any of the other nodes. Thus, $\{p, q, r\}$ is a path component, and similarly so is $\{s, t, u\}$. On the other hand, the second graph is path connected. Thus, by Lemma 2, the two graphs aren't isomorphic.

Before we move on, we'll add another reference object to our bag.

Definition. Let $n > 2$. The *cycle graph of order n* is

$$C_n = \{(v_1, \ldots v_n), \{v_1v_2, v_1v_n, v_2v_3, \ldots, v_{n-1}v_n\}).$$

A *cycle of length (n+1) in* Γ is a homomorphism $\sigma : C_n \to \Gamma$ which is injective on vertices.

Quick mental exercise: check that the two definitions of "cycle" are equivalent. Cycles are going to be very important in what is to come.

Embedded graphs

By now, those who are familiar with graphs are likely getting a little grumpy about the lack of pictures, and may have been serriptitiously drawing them while thinking about the examples. The reason we've been avoiding them is that drawing a picture of a graph isn't free: in doing so, we're assigning geometry to the purely combinatorial information in the graph, and that involves making choices. There will be a great many different pictures we can draw to represent a single graph, and each might influence our thinking about it in subtle ways. When dealing with data, we should always be acutely aware of what's actually given and what we've added, since any inferrences we make based on the latter are dangerous unless we understand exactly what we've chosen and how it affects the outcome.

Such geometry is often very useful, however. A "picture" of a graph, at least in the most familiar form, is an *embedding* (called a *realization*) of the graph in Euclidean space: a collection of points representing the vertices, and continuous arcs connecting them representing the edges.

Definition. Let $\Gamma = (V, E)$ be a graph. A *(geometric) realization* ρ of Γ in $\mathbb{R}^d$ is a function $\rho_V : V \to \mathbb{E}^d$ along with a collection of continuous arcs $\{\rho_{vv'} \subset \mathbb{E}^d\}_{vv' \in E}$ so that the endpoints of $\rho_{vv'}$ are $\rho_V(v)$ and $\rho_V(v')$ and arcs do not intersect themselves or other arcs except possibly at a shared endpoint.

This definition immediately suggests a question: can every graph be realized in every dimension?

It's quite easy to convince oneself that in $d = 0$ or $d = 1$, there are very few graphs with realizations – a good exercise is to characterize those that do.

On the other extreme, if $d = 3$ then every graph has a realization. To see this, take the vertices so that no three lie on a line, and no four lie in a plane, and use straight line segments for each edge. The former condition prevents edges from coinciding, while the latter keeps them from crossing at a point in their interiors. If either of these conditions doesn't hold for a given choice of vertices, just choose one of the vertices involved in the broken condition and perturb it slightly (in almost any direction). An arbitrarily small perturbation can always fix the condition without breaking any of the others, so we can always find a nearby collection of vertices on which to realize any given Γ. We say that sets vertices which satisfy all of these conditions are in *general position*. Above $d = 3$ the same game works.

That just leaves $d = 2$, the plane, where we can't play this trick: if we take the vertices of a convex quadrilateral, the diagonals cross and can't just be nudged out of doing so. There's hope, though, since we can take that diagonal move it elsewhere so the edges don't cross, or maybe move the vertices so these don't cross. Also, notice that adding edges can only make it harder to draw a proper embedding, so we can start by looking at graphs with all possible edges to find problems.

Definition. Let $V = \{v_1, v_2, \ldots v_n\}$. The *complete graph on n vertices* is $K_n = (V, \binom{V}{2})$.

It appears that we have the following:

Proposition 3. *K_5 does not admit a realization in $\mathbb{R}^2$.*

If it's true, there are graphs which cannot be embedded in $\mathbb{R}^2$. This being a math course, that calls for the creation of an adjective.

Definition. If there is a realization of a graph Γ in $\mathbb{R}^2$, we say Γ is *planar*.

How can we prove that K_5 isn't planar? Let's think a bit about the properties of such graphs to get a feel for them. Clearly we can restrict our attention to path connected graphs, since a graph with multiple components is planar only if all of its constituent pieces are.

The Euler characteristic

Now that we've made the decision to draw pictures, let's draw a bunch. We also got some mileage earlier out of counting things, so let's keep track of things we can count as we go.

A pattern pops out very quickly: if there are no "loops" in Γ, there is one more vertex than there are edges.

Definition. A graph Γ which is connected and contains no cycles is a *tree*.

Lemma 4. *Let Γ be a graph. If two vertices in Γ are connected by a path π, there is a simple path containing only edges from π which connects them.*

Proof. Exercise. □

Lemma 5. *Let Γ be a tree. Any two vertices in Γ are connected by a unique simple path.*

Proof. Exercise. □

Lemma 6. *If $\Gamma = (V, E)$ is a non-empty tree then $|V| - |E| = 1$.*

Proof. We proceed by strong induction on the number of vertices. Our inductive base will be the unique trees with one vertex and no edges, and two vertices and one edge, for which $|V| - |E| = 1$.

Now, assume that for $0 < k < n$, all trees with k vertices have $(k-1)$ edges. Suppose $\Gamma = (V, E)$ is a tree with n vertices. Select an edge $vv' \in E$. By Lemma 5, there is a unique simple path from v to v' in Γ which must be the edge vv'. Further, by Lemma 4, any other path connecting v to v' must include this edge, so the subgraph $\Gamma' = (V, E \setminus vv')$ has two path components C_1 and C_2, k_1 and k_2 vertices respectively, $k_1 + k_2 = n$. Because no subgraph of Γ contains a cycle, the induced subgraphs on C_1 and C_2 are each trees, so by our inductive hypothesis, C_1 has $(k_1 - 1)$ edges and C_2 has $(k_2 - 1)$ edges. Thus, Γ has $(k_1 - 1) + (k_2 - 1) + 1 = n - 1$ edges, as required. □

What about graphs that aren't trees? How do cycles interact with the plane? From the picture, it looks like they divide it up into pieces.

Theorem 7 (Jordan curve theorem)**.** *Let $\gamma : [0, 1] \to \mathbb{R}^2$ be a continuous function which is injective on $[0, 1)$ and with $\gamma(0) = \gamma(1)$, whose image is called a* Jordan curve*. The complement of the Jordan curve, $\mathbb{R}^2 \setminus \gamma([0, 1])$, consists of two path components: a bounded region and an unbounded region which share as their boundary the Jordan curve.*

Perhaps surprisingly, this is not trivial to prove, even if we had all the terminology[8]. Right now, we'll take it at face value – it really *looks* true – but we'll come around and prove a stronger version of it later in the course. Here, we rephrase it to obtain the following useful fact.

[8] What's a path component of a subset of $\mathbb{R}^2$? Just the same as in a graph: a maximial collection of points which can be pairwise joined by continuous paths. We'll give formal definitions once we've got the right terminiology.

Corollary 8. *Let Γ be a planar graph with a realization ρ in $\mathbb{R}^2$. The complement of the image under ρ of a cycle in Γ consists of two path components.*

Right now, there are different kinds of path components: bounded, and unbounded. This sort of heterogeneity always complicates terminology and forces us to consider special cases, so we like to remove them whenever possible. Today we're going to use a bit of geometry to line things up just so. If we were doing the same thing two weeks from now, we would choose a much cleaner (and more topological) *one-point compactification* to do the same job.

Definition (Stereographic projection)**.** Map $\mathbb{R}^2$ into $\mathbb{R}^3$ by $(x, y) \mapsto (x, y, 0)$, let $S^2 = \{(x, y, z) \in \mathbb{R}^3 \mid \sqrt{x^2 + y^2 + z^2} = 1\}$, and take $P = (0, 0, 1)$. Define the *stereographic projection* of a point $p \in S^2$ to be the point of intersection between the line through P and p and the (x, y)-plane.

This map is a bijection from the sphere minus the north pole to points in the plane. As such, it provides a way to make any realization of a graph in $\mathbb{R}^2$ into a realization in S^2. Any realization in S^2 can be perturbed to miss the north pole, so any graph realization in S^2 can be turned into a realization in $\mathbb{R}^2$. Thus, *2-spherical* graphs are exactly the same as planar graphs[9].

[9] Of course, it's hard to draw things on the sphere, so we'll also use this fact to allow us to work in the plane when drawing pictures, keeping in mind that we have a bit of extra flexibility when moving things around on the sphere.

Why do we care? When we move to the sphere, there are no more unbounded components – the unbounded component just becomes the region containing the north pole.

Definition. Let Γ be a connected planar graph and ρ a realization of Γ in S^2. The connected components of the complement of the spherical realization are called the *faces* of the realization, and the collection of all such is denoted F.

Now the relationship is clearer: we can add an edge either by introducing a new vertex or not. If we add a vertex, we don't change the number of faces. If we don't, we introduce a new Jordan curve and up the number of faces by one. Thus, there is a correspondence between (a choice of) a subset of edges in Γ and the faces in the realization.

Lemma 9. *Let $\Gamma = (V, E)$ be a connected graph. There is a tree $\tau(\Gamma) = (V, E' \subset E)$ which is a subgraph of Γ, called a* spanning tree *for Γ.*

Proof. Suppose Γ is a connected graph with $n > 0$ cycles. Select one, say $(vv_1, v_1v_2, \ldots, v_kv)$. Removing the edge vv_1 does not disconnect the graph, since the remaining edges of the cycle form an alternative path from v_1 to v, so any path that contained the edge can be altered to contain this sequence of edges instead. The subgraph $\Gamma' = (V, E \setminus vv_1)$ has fewer than n cycles, and so iterating this process a finite number of times must result in a tree, as required. □

This proof says that building a spanning tree precisely requires removing a set of edges that correspond to cycles. Cycles correspond to faces through Jordan curves, which cut existing faces in two pieces – and when we remove an edge that kills a cycle, we're "recombining" two faces. Let's put that together formally.

Definition. Given a planar graph Γ with realization ρ in S^2, choose a spanning tree $\tau(\Gamma) = (V, E')$ of Γ. Define the *spherical dual* graph $\tau^*(\Gamma, \rho) = (F, \widetilde{E})$ with vertices corresponding to the faces of the realization and an edge between two faces if they share a boundary edge that was removed in creating $\tau(\Gamma)$.

Lemma 10. *Let Γ be a planar graph with realization ρ in S^2 and spanning tree $\tau(\Gamma)$. The spherical dual $\tau^*_\rho(\Gamma)$ is a tree.*

Proof. Exercise. □

So, given a planar graph $\Gamma = (V, E)$, we can choose a realization ρ in S^2 and build a spanning tree $\tau(\Gamma) = (V, E')$ and its spherical dual tree $\tau^*(\Gamma, \rho) = (F, \widetilde{E} \cong E \setminus E')$. Lemma 6 then tells us that $|V| - |E'| = 1$ and that $|F| - |\widetilde{E}| = |F| - (|E| - |E'|) = 1$. Summing, we get

Theorem 11 (Euler characteristic of S^2)**.** *Let Γ be a connected planar graph with realization ρ in S^2. Then $|V| - |E| + |F| = 2$.*

That's very cool, but we were trying to prove that K_5 is non-planar, so we can't possibly be counting faces. We need to get rid of that F.

Corollary 12. *Let Γ be a connected planar graph with $|V| \geq 3$. Then $|E| \leq 3|V| - 6$.*

Proof. Exercise. □

Finally, we obtain

Corollary 13. *K_5 does not admit a realization in $\mathbb{R}^2$.*

Proof. K_5 has 5 vertices and $\binom{5}{2} = 10$ edges, and $10 \not\leq 3(5) - 6$. By Corollary 12, K_5 cannot be planar. □

Before we declare our mission accomplished, let's revisit our decision to map graphs to the sphere. What Theorem 11 appears to say is that any (nice) way we chop up the sphere into regions, we have some combinatorial relationship appearing. This association of the number 2 with the sphere is a very, very deep rabbit hole.

But we made a choice there – there are a lot of different shapes we can map these graphs onto. In particular, spherical polyhedra are the class of polyhedra that can be obtained by "flattening" the faces of a graph embedded on a sphere, so they all have this property. Another favorite of mine is the soccer ball – twelve pentagons and twenty hexagons. It's a weird coincidence that there are twelve pentagons on both of the objects with pentagons, isn't it?

Theorem 14. *Suppose $\mathcal{P}$ is a spherical polyhedron with only regular pentagonal and regular hexagonal faces. Then $\mathcal{P}$ has 12 pentagonal faces.*

Proof. Suppose there are P pentagons and H hexagons, so $|F| = P + H$. Each pentagon has five edges and vertices, and each hexagon has six. However, there are two faces sharing each edge, so we have $2|E| = 5P + 6H$. Finally, each vertex appears in three faces[10], so $3|V| = 5P + 6H$. Putting this all together, we get

$$2 = |V| - |E| + |F| = \frac{5}{3}P + 2H - \frac{5}{2}P - 3H + P + H = \frac{P}{6}.$$

So, $P = 12$. □

[10] Some geometry required. We don't *really* need the faces to be regular polyhedra, we just need three to meet at each vertex, which is guaranteed for the regular ones.

Topological spaces and continuous functions

OUR GOAL IS TO BUILD a general framework for understanding and comparing the intrinsic structure of objects. We will rely on intuition from Euclidean space as a roadmap.

Topological spaces

In $\mathbb{R}^d$, we (usually) measure distance using the *Euclidean metric*[11]. Points x and y are "close together" if the distance between them is small. How small? That depends on a notion of *scale*, which depends on context. As mathematicians, we like to use $\epsilon > 0$, so y is *ϵ-close* to x if $d(x, y) < \epsilon$ To conform with our plan to work with non-quantitative structures, though, what we really want is a check we can make without looking at distance directly.

[11] What is to follow will work in any metric space, though we don't have time to get into the details.

Definition. Let $\vec{x} \in \mathbb{R}^d$ and $\epsilon > 0$. The *open ϵ-ball around x* is

$$B_\epsilon(x) = \{y \mid d(x, y) < \epsilon\}.$$

Now, if someone hands us the collection of ϵ-balls around points we can ask, "Is $y \in B_\epsilon(x)$?" as a proxy for closeness without directly computing distance.

Another useful notion from Euclidean space is that of an *open set*.

Definition. Let $U \subseteq \mathbb{R}^d$. If for every $x \in U$ there is some $\epsilon(x) > 0$ with $B_{\epsilon(x)}(x) \subseteq U$, we call U an *open set*.

Thus, if we make an $\epsilon(x)$-small measurement error when attempting to select a point x from U, we still get a point from U – membership in such sets can be measured in a relatively robust manner. This is useful in a world of imprecise measurements and noise. Notice, we can flip this characterization around: an open set is one which is a union of open balls:

$$U = \bigcup_{x \in U} B_{\epsilon(x)}(x).$$

In particular, this is true *for open balls* – they are themselves open sets!

In our setting, we won't have all of these ϵs lying around, but we're going to mimic this notion.

Definition. Let X be a set. A *topology on* X is a collection $\tau \subseteq 2^X$ of *open sets* so that

i. $X \in \tau$ and $\varnothing \in \tau$,

ii. If $U_i \in \tau$ for all i in some index set A, then $(\bigcup_{i \in A} U_i) \in \tau$, and

iii. If $U_i \in \tau$ for all $i \in A$, $|A| < \infty$, then $(\bigcap_{i \in A} U_i) \in \tau$.[12]

A *topological space* is a pair (X, τ) where X is a set and τ is a topology on X. If $U \subseteq X$ has $U \in \tau$, we say U is *open in* X. Given a point $x \in X$, every open set containing x is an *open neighborhood of* x. A *closed set* is the complement of an open set; i.e. $C \subseteq X$ is closed if there is $U \in \tau$ so that $C = X \setminus U$.[13]

[12] We'd like to have arbitrary intersections of open sets, but sadly this won't work. Consider, for example, $\bigcap_{n=1}^{\infty} B_{\frac{1}{n}}(x) = \{x\}$, which is not an open set.

[13] Bad news! We have that $\varnothing$ is an open set, but X is open, so its complement, $\varnothing$, must be closed. Open and closed are *not opposites*!

Topologies are very complex objects, so we usually build them by specifying simpler families of sets that *generate* them, like our ϵ-balls from $\mathbb{R}^d$.

Definition. Let X be a set. If $\mathcal{B}$ is a collection of subsets of X, containing $\varnothing$ and X, so that every finite intersection of elements of $\mathcal{B}$ can be written as a union of elements of $\mathcal{B}$, then $\mathcal{B}$ is a *base* and *generates* some topology τ on X.

Examples.

- $\mathbb{R}^d$

  a. $\{B_\epsilon(x) \mid x \in X, \epsilon > 0\}$ is a base which generates the *standard* or *Euclidean topology* on X.

  b. $\{B_\epsilon(x) \mid x \in X, \epsilon > 0, \epsilon \in \mathbb{Q}\}$ also generates the standard topology.

  c. So does $\{B_\epsilon(x) \mid x \in \mathbb{Q}^k, \epsilon > 0, \epsilon \in \mathbb{Q}\}$

- X a set, $\tau = \{\varnothing, X\}$, the *trivial topology*.

- X a set, $\tau = 2^X$, the *discrete topology*.

- $X = \{p, q\}$, $\tau = \{\varnothing, \{p\}, \{p, q\}\}$, called the *Sierpinski space*

While it is occasionally convenient to describe a topological space from the ground up, it is usually conceptually cleaner and easier to understand a description if we describe it in relation to a space we already know.

Definition. Let (X, τ) be a topological space and $Y \subseteq X$. The subspace topology on Y is given by $\{U \cap Y \mid U \in \tau\}$.

Examples.

- Let $I = [0, 1] \subset \mathbb{R}$, the *standard closed interval*. What are the open sets in the subspace topology?
- The *standard d-sphere* is

$$S^d = \{x \in \mathbb{R}^{d+1} \mid ||x|| = 1\}$$

  Open sets in S^d under the subspace topology are given by intersections of Euclidean open sets with the sphere.

- The *standard d-disk* is

$$D^d = \{x \in \mathbb{R}^d \mid ||x|| \leq 1\}$$

  Observe that there is a canonical (identity) map $S^d \hookrightarrow D^{d+1}$. This is going to be one of our favorite maps.

- The *standard d-simplex* is the set

$$\Delta^d = \{x \in \mathbb{R}^{d+1} \mid \sum_{i=1}^{d+1} x_i = 1, x_i \geq 0, i = 1, \ldots, d+1\}.$$

  We should actually be thinking of the d-simplex as a subspace of the hyperplane $\sum_{i=1}^{d+1} x_i = 1$ which is itself a subspace of $\mathbb{R}^{d+1.}$

Continuous functions

Just like with graphs, we want a notion of how to compare topological spaces. The "correct" choice of maps should preserve the data of the topological space: open sets.

We might want to say that $f : X \to Y$ is continuous if it sends open sets to open sets. However, consider the constant map

$$f : \mathbb{R} \to \mathbb{R}, \qquad x \mapsto 0.$$

Clearly this should be continuous, but it maps the open set $\mathbb{R}$ to the non-open set $\{0\}$.

Note that a function f is continuous at x if the image of a small neighborhood N_x of x, $f(N_x)$ gets arbitrarily small as N_x does. That is, if we consider a small open neighborhood of $f(x)$, it should be the image of a small open neighborhood of x.

Definition. Let X and Y be topological spaces. A function $f : X \to Y$ is *continuous* if for every open set $U \subseteq Y$, the set $f^{-1}(U) \subseteq X$ is open. We call continuous functions *maps*.

In our constant function example above, if $U \subseteq \mathbb{R}$ is an open neighborhood of zero, $f^{-1}(U) = f^{-1}(0) = \mathbb{R}$ which is open. If V is an open set in $\mathbb{R}$ which doesn't contain zero, then $f^{-1}(V) = f^{-1}(\varnothing) = \varnothing$, which is also open. Thus, the function is continuous.

Theorem 15. *For a function $f : \mathbb{R} \to \mathbb{R}$, the above definition of continuity agrees with the usual ϵ-δ definition from calculus.*

Proof. Exercise. □

Our definitions of the subspace topology is motivated by the desire to make the natural inclusion of sets into a map in the simplest possible way.

Lemma 16. *Let X be a topological space, $Y \subseteq X$ equipped with the subpace topology. The* inclusion map $\iota : Y \hookrightarrow X$ *given by $\iota(x) = x$ is continuous.*

Proof. Let $U \subseteq X$ be open. The preimage is $\iota^{-1}(U) = \{y \in U \mid y \in Y\} = U \cap Y$, which is open in Y by definition, so ι is continuous. □

One of the most powerful tools we'll have in our quest to build new topological spaces are *quotient spaces*. Quotients provide us with the power to "glue pieces together" or "forget information", depending on context. We build quotient spaces by defining an *equivalence relation*[14] on the space X and then building the simplest topology that makes the induced *projection* map continuous.

[14] Recall that an equivalence relation on a set X is a relation $\sim$ which is reflexive, symmetric and transitive. Write $[x]$ for the *equivalence class* of x. The equivlance classes under $\sim$ partition X, so we can think of the collection of classes as a new set $(X/\sim) = \{[x] \mid x \in X\}$, the *quotient of X by $\sim$*.

Definition. Let (X, τ) be a topological space and $\sim$ an equivalence relation on X. The quotient map $\pi : X \to (X/\sim)$ is given by $\pi(x) = [x]$, and the quotient topology on $(X/\sim)$ is such that $U \in (X/\sim)$ is open if and only if the preimage $\pi^{-1}(U)$ is open.

Quotient spaces are one of the most powerful tools we have for building spaces.

Examples.

- Let $\Gamma = (V, E)$ be a graph. The *topological realization* of Γ is the topological space given by[15]

$$\Big(\coprod_{v \in V} \Delta^0 \times \coprod_{vv' \in E} \Delta^1\Big) / ((1,0)_{vv'} \sim v, (0,1)_{vv'} \sim v').$$

  Open sets in this space inside of edges are just the same as usual. If they cross a vertex, they've got to bleed into *all* nearby edges.

[15] This is a lie. In order to do this correctly, what we really need to do is make a choice for each edge of a vertex to attach at 0 and the other to attach at 1. That requires notation that obscures a simple idea, so I didn't write it down.