APPLIED TOPOLOGY LECTURE NOTES

CHAD GIUSTI

UPDATED: SEPTEMBER 1, 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 $v_1v_2 \in E$ and $v_2v_3 \in E$ 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]

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]

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 $\widetilde{\varphi}(vv') = \varphi(v)\varphi(v')$.

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

[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.

[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.

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

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 wy, we must have $\varphi(q) = x$ or y and $\varphi(s)$ as the other. Either choice forces $\varphi(r) = z$, 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, . . . v_n\, \v_1v_2, v_2v_3, . . . , v$_{n-1}$v_n\.

Definition. Let $\Gamma = (V, E)$ be a graph with $a, b \in V$. A *path of length n 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 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* $ß = 1, \ldots, n$*, 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 *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.