Math 7H	Professor: Padraic Bartlett

Lecture 4: Induction, the Euler Characteristic, and Chemistry

Week 4 — *UCSB 2015*

Today's lecture is a strange one; in it, we're going to use a fundamental proof technique (induction) to prove a deep and beautiful fact in topology (the Euler characteristic) to answer a problem in chemistry! It might seem sort of strange, but in fact these sorts of connections are what mathematics becomes as you go further into it; lectures are almost never "let's talk about this one idea," but rather become "what happens when you stick together ideas x,y and z?"

Also, it's gorgeous! We start with the fundamental idea for today: induction!

1 Induction

Sometimes, in mathematics, we will want to prove the truth of some statement $P(n)$ that depends on some variable n. For example:

- $P(n) =$ "The sum of the first n natural numbers is $\frac{n(n+1)}{2}$."
- $P(n) =$ "If $q \geq 2$, we have $n \leq q^n$.
- $P(n) =$ "Every polynomial of degree n has at most n roots."

For any fixed n, we can usually use our previously-established methods to prove the truth or falsity of the statement. However, sometimes we will want to prove that one of these statements holds for **every** value $n \in \mathbb{N}$. How can we do this?

One method for proving such claims for every $n \in \mathbb{N}$ is the method of **mathematical induction**! Proofs by induction are somewhat more complicated than the previous two methods. We sketch their structure below:

- To start, we take our claim $P(n)$, that we want to prove holds for every $n \in \mathbb{N}$.
- The first step in our proof is the **base step**:in this step, we explicitly prove that the statement $P(1)$ holds, using normal proof methods.
- With this done, we move to the **induction step** of our proof: here, we prove the statement $P(n) \implies P(n+1)$, for every $n \in \mathbb{N}$. This is an implication; we will usually prove it directly by assuming that $P(n)$ holds and using this to conclude that $P(n+1)$ holds.

Once we've done these two steps, the principle of induction says that we've actually proven our claim for all $n \in \mathbb{N}$! The rigorous reason for this is the **well-ordering principle**, which we discussed in class; however, there are perhaps more intuitive ways to think about induction as well.

The way I usually think of inductive proofs is to think of toppling dominoes. Specifically, think of each of your $P(n)$ propositions as individual dominoes – one labeled $P(1)$, one labeled $P(2)$, one labeled $P(3)$, and so on/so forth. With our inductive step, we are insuring that all of our dominoes are *lined up* – in other words, that if one of them is true, that it will "knock over" whichever one comes after it and force it to be true as well! Then, we can think of the base step as "knocking over" the first domino; once we do that, the inductive step makes it so that all of the later dominoes also have to fall, and therefore that our proposition must be true for all n (because all the dominoes fell!)

To illustrate how these kinds of proofs go, here's an example:

Claim. For any $n \in \mathbb{N}$, take a $2^n \times 2^n$ grid of unit squares, and remove one square from somewhere in your grid. The resulting grid can be tiled by 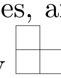 - shapes.

Proof. As suggested by the section title, we proceed by induction.

Base case: for $n = 1$, we simply have a 2×2 grid with one square punched out. As this *is* one of our three-square shapes, we are trivially done here.

Inductive step: Assume that we can do this for a $2^k \times 2^k$-grid without a square, for any $k \leq n$. We then want to prove that we can do this for a $2^{n+1} \times 2^{n+1}$ grid minus a square.

So: take any such grid, and divide it along the dashed indicated lines into four $2^n \times 2^n$ grids. By rotating our grid, make it so that the one missing square is in the upper-right hand corner, as shown below:

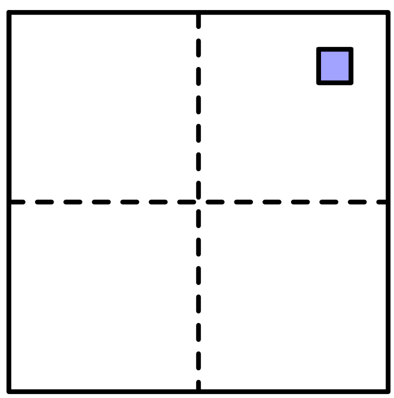

Take this grid, and carefully place down one thre-square shape as depicted in the picture below:

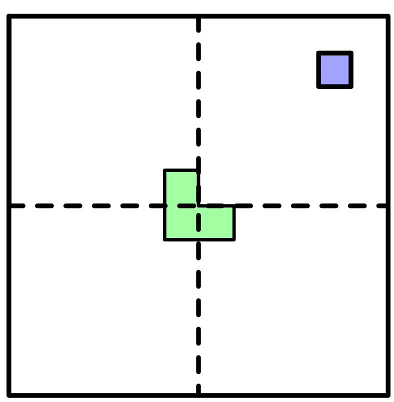

Now, look at each of the four $2^n \times 2^n$ squares in the above picture. They all are missing exactly one square: the upper-right hand one because of our original setup, and the other three because of our placed three-square-shape. Thus, by our inductive hypothesis, we know that all of these squares can also be tiled! Doing so then gives us a tiling of the whole shape; so we've created a tiling of the $2^{n+1} \times 2^{n+1}$ grid!

As this completes our inductive step, we are thus done with our proof by induction. $\square$

Let's look at another example of an inductive problem:

Example. Draw some straight lines in the plane. Notice that when we do this, we divide the plane up into regions bounded by these lines. What is the maximum number of regions we can divide the plane into with n lines?

Answer. Again, take a moment to work out some base cases and figure out what's going on here!

Here's a few observations you're likely to have made:

1. No lines break the plane into one piece, as we've not split anything up! One line breaks the plane into two pieces; two lines breaks the plane into up to four pieces if those lines are not parallel; three lines can break the plane up into seven pieces if we are careful to not let all three lines intersect at the same place; and four lines can break the plane up into up to 11 pieces if we are again careful to not have any more than two lines intersect at any point, and also not have any parallel lines! In general, it looks like n lines is giving us $1 + \frac{n(n+1)}{2}$ regions, given enough data and staring at things.

2. In general, it looks like the n-th line is adding at most n new regions to our plane. To see why this might hold in general, consider the process of drawing any line.

   (a) If our line intersects any region, it divides that region into two pieces! This is the only way our line creates new regions.

   (b) Our line enters a region if and only if it crosses one of the lines that bounds that region.

   (c) As well, before our line crosses any other regions, it by default starts in some region already.

   (d) Therefore, the total number of times our line intersects other lines, plus one, is the total number of new regions created!

   (e) We can cross each other line at most once, as our lines are straight.

   (f) Therefore, if there are n lines in existence, we can create at most $n + 1$ new regions by adding a $n + 1$-th line.

3. Furthermore, notice that it is always possible to draw such a line! To draw a line, we need to give two pieces of information:

   (a) Its slope needs to not be parallel to any other existing line's slope, to insure that it can intersect that line. There are only n slopes currently used and infinitely many possibilities, so this is always possible.

(b) Given a slope, we need to pick a x-intercept for our line. Furthermore, we want to do this so that our line does not intersect any other lines at places where multiple lines are already intersecting: this would make our line use up multiple "intersecting other line" instances, while only entering one region (which means we wouldn't get to $n+1$!) There are only finitely many such existing intersection points, and infinitely many choices of x-intercept; so we can also avoid all of these possibilities.

Therefore, it is possible to always draw a line that intersects other lines in n places, and thus that creates $n+1$ new regions!

By the above, we have a rather nice recurrence relation: if L_n is the total number of regions that we can divide the plane up into with n lines, we have

$$L_{n+1} = (n+1) + L_n.$$

Using this, we can prove that our guess of $L_n = 1 + \frac{n(n+1)}{2}$ is right via induction:

Base case: We know $L_0 = 1 = 1 + \frac{0 \cdot 1}{2}$ from our case work.

Ind. step: Assume that $L_n = 1 + \frac{n(n+1)}{2}$ for each n from 1 to m; we will seek to prove that $L_{m+1} = 1 + \frac{(m+1)(m+2)}{2}$. This is pretty quick: notice that

$$\begin{aligned} L_{m+1} &= (m+1) + L_m, \text{ by our recurrence relation,} \\ &= (m+1) + \frac{m(m+1)}{2} + 1, \text{ by our inductive hypothesis,} \\ &= \frac{m^2 + m + 2m + 2}{2} + 1 \\ &= \frac{(m+1)(m+2)}{2} + 1, \text{ as claimed.} \end{aligned}$$

2 Euler characteristic

To start this part of the talk, I need to introduce a very useful concept that we'll refer to in many future lectures: **graphs**!

Definition. A **graph** G with n vertices and m edges consists of the following two objects:

1. a set $V = \{v_1, \dots v_n\}$, the members of which we call G's **vertices**, and

2. a set $E = \{e_1, \dots e_m\}$, the members of which we call G's **edges**, where each edge e_i is an unordered pair of distinct elements in V, and no unordered pair is repeated. For a given edge $e = \{v, w\}$, we will often refer to the two vertices v, w contained by e as its endpoints.

Example. The following pair (V, E) defines a graph G on five vertices and five edges:

- $V = \{1, 2, 3, 4, 5\}$,

- $E = \{\{1,2\},\{2,3\},\{3,4\},\{4,5\},\{5,1\}\}$.

Something mathematicians like to do to quickly represent graphs is **draw** them, which we can do by taking each vertex and assigning it a point in the plane, and taking each edge and drawing a curve between the two vertices represented by that edge. For example, one way to draw our graph G is the following:

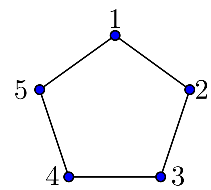

We could also draw our graph like this:

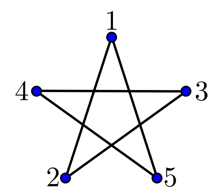

One nice property of a graph is when we can draw it in such a way that none of the edges cross. We call this property **planarity**:

Definition. We say that a graph G is **planar** if we can draw it in the plane so that none of its edges intersect.

For example, the pentagon above is planar, because you can draw it without having any of its edges cross with the first drawing! (In particular, this graph is planar even though the second drawing **does** have crossing edges; all we need is that there is **some** way to draw it without crossings.)

One useful notion for a planar graph is the idea of "faces," i.e. the regions bounded by the edges of our graph.

Definition. For any planar graph G, we can define a **face** of G to be a connected region of $\mathbb{R}$ whose boundary is given by the edges of G.

For example, the following graph has four faces, as labeled:

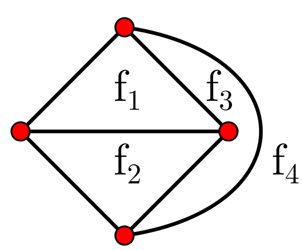

Planar graphs have a surprising property: for any planar graph G, if we let V denote the number of edges in the graph, E denote the number of edges, and F denote the number of faces, then $V - E + F$ is always the same! In particular, it is always 2! We prove this fact, called the **Euler characterstic** of the plane:

Theorem. (Euler characteristic.) Take any graph that has been drawn in $\mathbb{R}^2$ as a planar graph. Then, if V is the number of vertices, E is the number of edges, and F is the number of faces in this graph, we have the following relation:

$$V - E + F = 2.$$

Proof. We will actually prove a stronger claim: we will show that any planar **multigraph** (a graph, but where we allow multiple edges between vertices, and also edges that start and end at the same vertex) satisfies the $V - E + F = 2$ formula. For the rest of this proof, we will assume that graph and multigraph are synonymous; once we are done with this proof, though, we will stop assuming this.

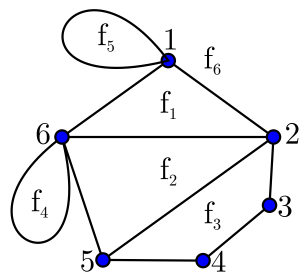

We proceed by induction on the number of vertices. Suppose that $V = 1$. Then our graph looks like something of the following form:

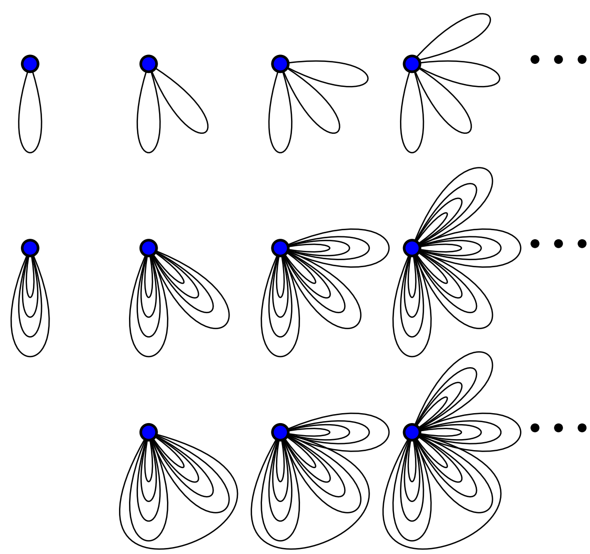

I claim that $V - E + F = 2$ for any of these graphs, and prove it by a second induction on the number of edges. For a zero-edge graph, this is easy; we have one vertex, no edges and one face, we have $V - E + F = 1 - 0 + 1 = 2$. Now, assume via induction that every one-vertex multigraph on n edges has $V - E + F = 2$. Take any graph on one vertex with $n + 1$ edges. Pick one of these edges, and look at it.

I claim that this edge borders exactly two faces. To see why, take any edge, and assign an orientation to it (i.e. if our edge is $\{x, y\}$, then orient the edge so that we travel from

x to y.) If you do this, then our edge has two "sides," the left- and right-hand sides, if we travel along it via this orientation.

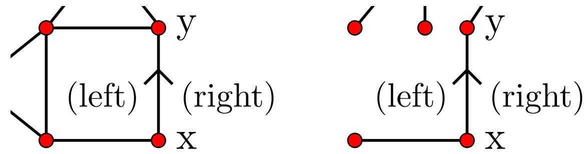

There are two possibilities, as drawn above: either the left- and right-hand sides are different, or they are the same. This tells us that our edge either borders one or two faces! To see that we have exactly two, we now recall that our edge (because our graph has exactly one vertex) must start and end at the same vertex. In other words, it is a **closed** loop: i.e. its outside is different from its inside! In other words, our left- and right-hand sides are different, and our edge separates two distinct faces.

Therefore, deleting this edge does the following things to the graph: it decreases our edge count by 1, and also decreases our face count by 1 (as we merge two faces when we delete this edge.) In other words, deleting this edge does not change $V - E + F$! But by induction we know that $V - E + F = 2$ for all 1-vertex graphs on n edges, which is what we get if we delete this edge from a $n + 1$-edge graph. So we're done!

This settles our base case for our larger induction on V, the number of vertices. We now go to the second phase of an inductive proof: we show how to reduce larger cases to smaller cases!

To do this, consider the following operation, called **edge contraction**. Take any edge with two distinct endpoints. Delete this edge, and combine its two endpoints together: this gives us a new graph! We draw examples of this process below: we start with a graph on six vertices, and contract one by one the edges labeled in red at each step.

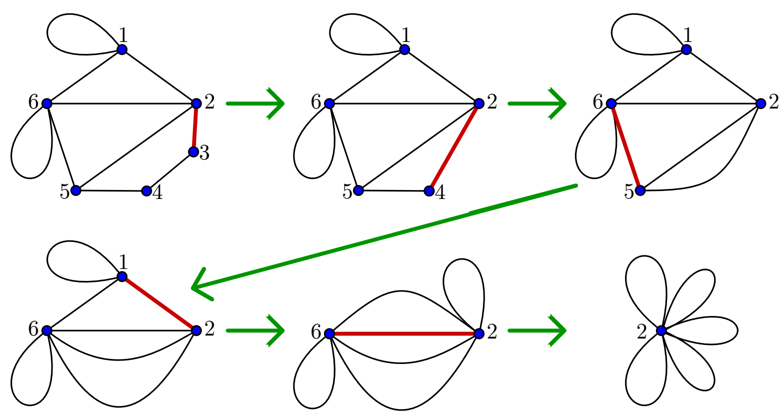

Contracting an edge decreases the number of vertices by 1 at each step, as it "squishes together" two adjacent vertices into one vertex. It also decreases the number of edges by 1 at each step, as we are contracting an edge to a point! Finally, it never changes the number of faces; if two faces were distinct before this process happens, they stay distinct, as we're not making any cuts in any of our boundaries (and instead are just shrinking them partially a bit!)

But this means that $V - E + F$ is still constant! Therefore, by induction, if $V - E + F$ holds for every n-vertex multigraph, it holds for any $n+1$-vertex multigraph by just contracting an edge! This finishes our induction, and thus our proof. $\square$

3 Chemistry

In chemistry, a **fullerene**[1] is (roughly) any molecule made entirely out of carbons that makes a sphere or ellipsoid. Fullerenes have lots of strange applications ranging from cancer treatments to superstrong materials; some people even conjecture that they are the seed behind all organic life on Earth.

Accordingly, we'd like to understand them! To do this, let's use a bit of mathematics. In chemistry, carbon molecules want to do very specific things:

- Each carbon wants to be connected to three other carbons.
- Those connections do not want to cross.
- The faces formed by this graph want to all be either 5-cycles[2] and 6-cycles. (Cycles that are smaller or larger make the carbon molecules unstable and prone to breaking.)

With these simple restrictions, you can prove the following remarkable result:

Proposition 1. *Any fullerene has precisely 12 5-cycles.*

Proof. This is this week's HW! $\square$

In chemistry, it seems that not every fullerene is realizable. In particular, one rule that chemists have noticed that all fullerenes obey is that **they never have two adjacent pentagonal faces**: this is probably because the pentagon is not a shape that carbons are terribly happy in, and the stress of having any carbon in two such faces probably makes any such molecule unstable.

Therefore, it seems likely that any viable fullerene will have to have all of its pentagonal faces isolated. By the proposition above, it must have at least 60 vertices, as it has precisely 12 such faces, and each face has 5 vertices. So: does such a fullerene exist?

The answer, as it turns out, is yes! We draw it stretched out as a planar graph on the next page:

---

[1] In popular science fiction, a **fullerene** is a plot device that basically allows you to do whatever you want. See also: wormholes, quantum mechanics, hot cups of tea.

[2] The **cycle graph** on n vertices, C_n, is the graph on the vertex set $\{v_1, v_2, \ldots v_n\}$ with edge set $E(C_n) = \{\{v_1, v_2\}, \{v_2, v_3\}, \ldots \{v_{n-1}, v_n\}, \{v_n, v_1\}\}$. The cycle graphs C_n can be drawn as n-gons, as depicted below:

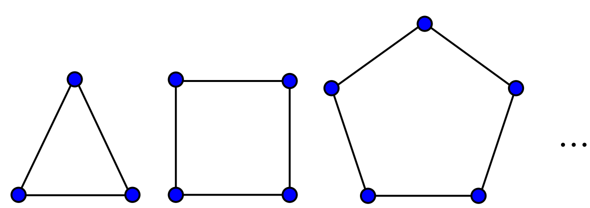

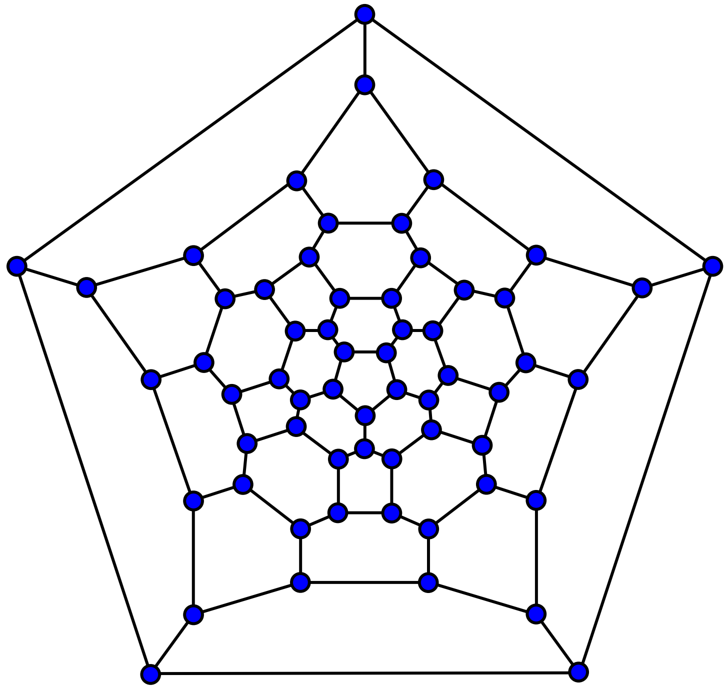

And also as a sphere, which is what it's like when we don't stretch it out:

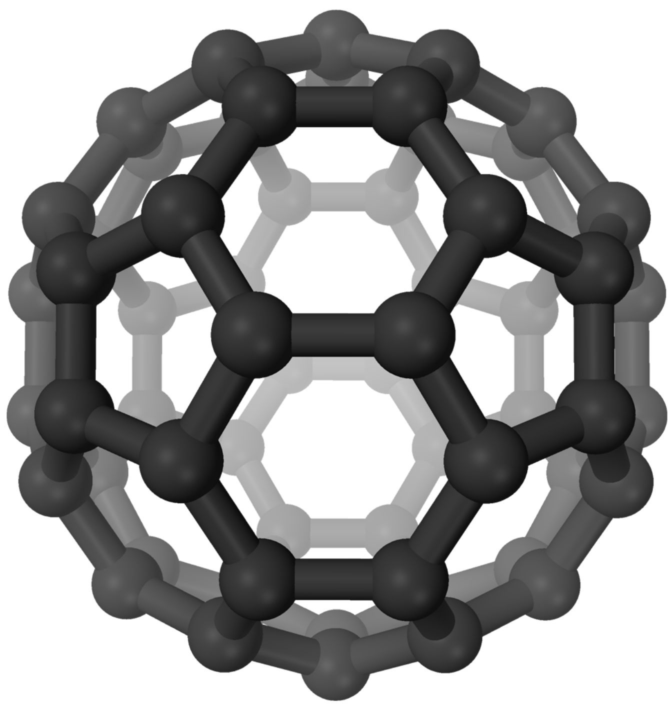

Cool, right?