Lecture 5: Infinite Random Graphs

In today's talk, we are going to discuss the concept of **random graphs**! Specifically, consider the following definition:

Definition. Take n vertices $\{v_1, \ldots v_n\}$. For each pair of vertices v_i, v_j, flip a coin that comes up heads half of the time and tails the other half of the time, and connect v_i to v_j with an edge if and only if our coin came up heads.

We call the result of this process a **random graph on n vertices with respect to the uniform distribution,** and denote this as $G_{n,1/2}$ for shorthand.

Notice that in this process there are $\binom{n}{2}$ possible edges, and for each edge it exists or it doesn't: so in the end we have $2^{\binom{n}{2}}$-many possible graphs. Furthermore, note that this process thinks that any specific graph on $\{v_1, \ldots v_n\}$ is 'equally likely" to come up: i.e. given any specific graph H on these n vertices, we need to have every coin flip come up a specific way in $G_{n,1/2}$ to get H. This is why we have the word "uniform" here: it's because our process is as likely to pick out one graph as any other.

Given a random graph on n vertices, we can ask lots of questions about it, like the following:

As a warmup, we start with the following question:

Question 1. *Let $e(H)$ denote the function that takes in a graph H and outputs the total number of edges in H. What is the expected value of e over $G_{n,1/2}$? In other words, if you take a random graph on n vertices under our model, how many edges would you expect to see on average?*

Answer. If you think of the model we have for our random graph, this is pretty clear: if you have n vertices and you're flipping a coin for each pair of them, you'd expect to see $\frac{\binom{n}{2}}{2}$ many heads, i.e. $\frac{\binom{n}{2}}{2}$ many edges.

More interestingly, we can ask how likely our random graph is to contain a given **structure**. For example, instead of just asking how many edges our graph has, we could ask how many distinct labeled triangles occur as subgraphs of our graph:

Question 2. *Let $t(H)$ denote the function that takes in a graph H and outputs the total number of distinct labellings of triangles in H. What is the expected value of t over $G_{n,1/2}$? In other words, if you take a random graph on n vertices under our model, how many triangles would you expect to see on average?*

Answer. We calculate the expected value of this function t as follows:

- Take a random graph on n vertices.

- Pick any three vertices a, b, c in this graph.
- Find the probability that these three specific vertices make a triangle.
- Now, multiply this probability by the number of three-vertex sets: this gives us the total expected number of triangles in our graph!

Under this model, the total number of triangles we would expect to see is

$$\begin{aligned} &\sum_{a,b,c \text{ distinct labeled vertices}} \mathbb{P}(abc \text{ is a triangle}) \\ =&\sum_{a,b,c \text{ distinct labeled vertices}} \mathbb{P}(ab \text{ is an edge}) \cdot \mathbb{P}(bc \text{ is an edge}) \cdot \mathbb{P}(ac \text{ is an edge}) \\ =&\sum_{a,b,c \text{ distinct labeled vertices}} \frac{1}{8} \\ =&n \cdot (n-1) \cdot (n-2) \cdot \frac{1}{8}, \end{aligned}$$

because there are $n(n-1)(n-2)$ many ways to choose three vertices along with a labeling, and the probability that that individual triple of edges forms a triangle is $\frac{1}{8}$.

For $n = 3$, this formula predicts that we'll see $\frac{3!}{0!} \cdot \frac{1}{8} = \frac{3}{4}$ triangles on average. However, if we look at the collection of all graphs on three vertices, we see that there are eight possible graphs, only one of which is a triangle; so we'd expect the probability here to be $\frac{1}{8}$. Why is there this incongruity between our answers?

Well, because in our problem we're counting the number of **distinct labellings of triangles**, not just triangles! To be precise, we are counting all of the ways of mapping the three numbers $1, 2, 3$ onto a set of three vertices in a random graph, so that the resulting image under this mapping is a triangle. Therefore, whenever there is a triangle in our random graph, we actually wind up counting it six times:

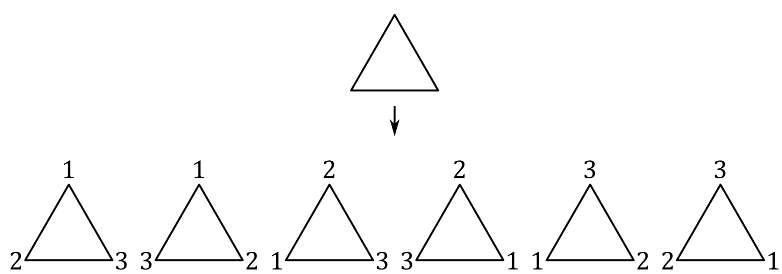

This accounts for the discrepancy between our two counts: while we might have expected $\frac{1}{8}$, we will really get $\frac{6}{8}$ if we want to count distinct labeled occurrences of our triangle.

If we want to ignore this question of distinct labellings, and just count the number of different triangles, we can just divide our formula by the number of symmetries of a triangle, because this is the same thing as the number of ways of labeling a triangle and having it remain a triangle. In other words, a formula for counting distinct triangles without caring about labellings is

$$n \cdot (n-1) \cdot (n-2) \cdot \frac{1}{6} \cdot \frac{1}{8}.$$

For $n = 4$, this predicts that your typical graph will contain about half a triangle. If you enumerate all of the 64 graphs on 4 vertices and count, you can see that there are 32 triangles, and therefore that this prediction aligns with reality there.

1 The Rado Property

: In today's talk, we're going to focus on the following property of random graphs, and ask when a random graph has this property:

Definition. Let (‡) denote the **Rado** property of graphs: we say that a graph G satisfies the property (‡) iff for any pair of finite disjoint subsets $U, W \subset V(G)$, there is some $v \in V(G)$, $v \notin U \cup W$, such that v has an edge to every vertex in U and to no vertices in W.

Notice that no **finite** graph has this property, because we could just make $U =$ all of G, and we would be unable to find a $v \notin U$. But what if we considered an **infinite** random graph?

Definition. Take $\mathbb{N}$ vertices $\{v_i\}_{i=1}^{\infty}$. For each pair of vertices v_i, v_j, flip a coin that comes up heads half of the time and tails the other half of the time, and connect v_i to v_j with an edge if and only if our coin came up heads.

We call the result of this process a **random graph on $\mathbb{N}$ vertices with respect to the uniform distribution,** and denote this as $G_{\mathbb{N},1/2}$ for shorthand.

Given this definition, we have the following rather remarkable result:

Theorem 3. *If G is a random graph of the form $G_{\mathbb{N},p}$, for $p \neq 0, 1$, then G satisfies (‡) with probability 1.*

Proof. Choose any pair of finite disjoint subsets U, W in $V(G)$. Then, for any vertex $v \in V(G), v \notin U \cup W$, let A_v be the event that v is connected to all of U and none of W. Then, we have that

$$Pr(A_v) = p^{|U|} \cdot (1-p)^{|V|} > 0.$$

Because the probability that A_v doesn't happen plus the probability that A_v **does** happen must sum to 1, we then know that

$$Pr(\text{not } A_v) = 1 - p^{|U|} \cdot (1-p)^{|V|} = \lambda < 1,$$

for some constant $\lambda \in (0, 1)$.

Thus, we know that the probability of k different vertices $v_1, \ldots v_k$ all failing to satisfy A_v is λ^k, which goes to 0 as k increases! So we can specifically bound this probability above by ϵ, for any $\epsilon > 0$, by simply looking at enough vertices.

Now, note that there are only countably many pairs of finite disjoint subsets of $\mathbb{N}$; consequently, we can enumerate all such pairs in a list $\{(U_i, W_i)\}_{i=1}^{\infty}$, and bound the probability of (U_i, W_i) failing to have a vertex that hits all of U_i and none of W_i by $\epsilon/2^i$, for every i. Then, the probability of any of these events failing is bounded above by the sum

$$\sum_{n=1}^{\infty} \frac{\epsilon}{2^n} = \epsilon;$$

so thie probability of our graph satisfying property (‡) is greater than $1 - \epsilon$, for any $\epsilon > 0$; i.e. the probability of our graph satisfying this property is 1! So, almost every random graph satisfies property (‡). □

So: the probabilistic method is a fantastically useful way to show the existence of graphs with certain properties! However, it's not so great for actually providing concrete examples of such graphs; typically, an application of probabilistic ideas will only tell you that most graphs have your property, not what one such graph might actually look like.

In the light of the above comments, it's interesting to note that we can actually construct a graph that satisfies (‡)! In fact, consider the following construction:

Definition. The **Rado** graph is the following graph:

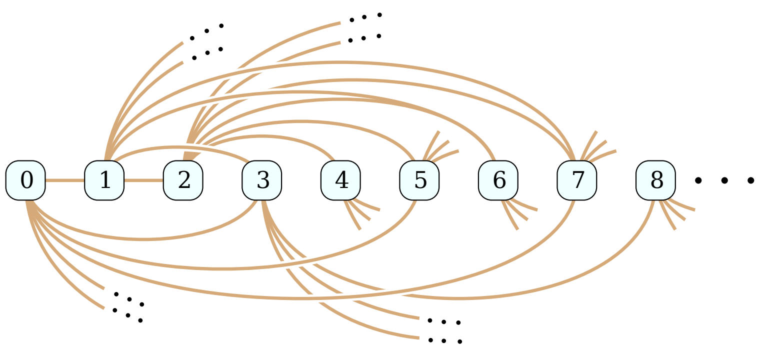

(Shamelessly stolen from Wikipedia.)

- Start by defining $R_0 = K_1$, the graph with a single vertex.
- If R_k is defined, let R_{k+1} be defined by the following: take R_k, and add a new vertex v_U for every possible subset U of R_k's vertices. Now, add an edge from v_U to every element in U, and to no other vertices in R_k.
- Let $R = \cup_{k=1}^{\infty} R_k$.

We claim that this is a graph on $\mathbb{N}$-many vertices that satisfies property (‡). To see why: pick any two finite disjoint subsets U, V of $V(R)$. Because each vertex of R lives in some R_k, we know that there is some value n such that U, V are both in fact subsets of R_n, as there are only finitely many elements in $U \cup V$. Then, by construction, we know that there is a vertex v_U in R_{n+1} with an edge to every vertex in U and to none in V.

We close our study of (‡) with the following proposition:

Proposition 4. *Any two graphs that satisfy (‡) are isomorphic*[1]*.*

Proof. To see this, take any two graphs $G = (V, E), G' = (V', E')$ on $\mathbb{N}$-many vertices that satisfy (‡); we will exhibit an isomorphism $\phi" V \to V'$ between them.

To do this: fix some ordering $\{v_i\}_{i=1}^{\infty}$ of V's vertices, and do the same for V'. We start with ϕ undefined for any values of V, and construct ϕ via the following back-and-forth process:

- At odd steps:
  - Let v be the first vertex under V's ordering that we haven't defined ϕ on, and
  - let U be the collection of all of v's neighbors in V that we currently **have** defined ϕ on.
  - By (‡), we know that there is a $v' \in V'$ such that v' is adjacent to all of the vertices in $\phi(U)$ and no other yet-defined vertices in V' that ϕ hits yet (as both sets are stil finite.)
- At even steps: do the exact same thing as above, except switch V and V'.

So, in other words, we're starting with ϕ totally undefined; at our first step, we're then just taking ϕ and saying that it maps $v_1 \in V$ to some element in V'. Then, at our second step, we're taking the smallest element in V' that's not $\phi(v_1)$, and mapping it to some element w that either does or does not share an edge with v, depending on whether $\phi(w)$ and $\phi(v)$ share an edge.

By repeating this process, we eventually get a map that's defined on all of V, V'; we claim that such a map is an isomorphism. Proving this is on the HW! ☐

Combining our results gives us the following rather surprising result:

Corollary 5. *With probability 1, any two random graphs are isomorphic.*

(... wait, what?)

[1]We say that two graphs G_1, G_2 are **isomorphic** if there is a bijection $\varphi : V(G_1) \to V(G_2)$, such that $\{x, y\}$ is an edge in G_1 if and only if $\{\varphi(x), \varphi(y)\}$ is an edge in G_2. In other words, there is a way to "relabel" the vertices of G_1 so that it looks like G_2: i.e. G_1 and G_2 are the "same," up to isomorphism.