Theorem 4.5 Let H be *any* fixed graph. For every subgraph H' of H (including H itself) let $X_{H'}$ denote the number of copies of H' in $G(n,p)$. Assume p is such that $E[X_{H'}] \rightarrow \infty$ for every H'. Then

$$X_H \sim E[X_H]$$

almost always.

Proof. Let H have v vertices and e edges. As in Theorem 4.4 it suffices to show $\Delta^* = o(E[X])$. We split Δ^* into a finite number of terms. For each H' with w vertice and f edges we have those $(y_1, \ldots, y_v)$ that overlap with the fixed $(x_1, \ldots, x_v)$ in a copy of H'. These terms contribute, up to constants,

$$n^{v-w}p^{e-f} = \Theta\left(\frac{E[X_H]}{E[X_{H'}]}\right) = o(E[X_H])$$

to Δ^*. Hence Corollary 3.5 does apply. $\Box$

be that subgraph with maximal density $\rho(H_1) = e_1/v_1$. (When H is balanced we may take $H_1 = H$.) They showed that $p = n^{-v_1/e_1}$ is the threshold function. This will follow fairly quickly from the methods of theorem 4.5.

We finish this section with two strengthenings of Theorem 4.2.

Theorem 4.4 Let H be strictly balanced with v vertices, e edges and a automorphisms. Let X be the number of copies of H in $G(n,p)$. Assume $p \gg n^{-v/e}$. Then almost always

$$X \sim \frac{n^v p^e}{a}$$

Proof. Label the vertices of H by $1,\ldots,v$. For each ordered $x_1,\ldots,x_v$ let $A_{x_1,\ldots,x_v}$ be the event that $x_1,\ldots,x_v$ provides a copy of H in that order. Specifically we define

$$A_{x_1,\ldots,x_v} : \{i,j\} \in E(H) \Rightarrow \{x_i,x_j\} \in E(G)$$

We let $I_{x_1,\ldots,x_v}$ be the corresponding indicator random variable. We define an equivalence class on v-tuples by setting $(x_1,\ldots,x_v) \equiv (y_1,\ldots,y_v)$ if there is an automorphism σ of $V(H)$ so that $y_{\sigma(i)} = x_i$ for $1 \le i \le v$.Then

$$X = \sum I_{x_1,\ldots,x_v}$$

gives the number of copies of H in G where the sum is taken over one entry from each equivalence class. As there are $(n)_v/a$ terms

$$E[X] = \frac{(n)_v}{a} E[I_{x_1,\ldots,x_v}] = \frac{(n)_v p^e}{a} \sim \frac{n^v p^e}{a}$$

Our assumption $p \gg n^{-v/e}$ implies $E[X] \to \infty$. It suffices therefore to show $\Delta^* = o(E[X])$. Fixing $x_1,\ldots,x_v$,

$$\Delta^* = \sum_{(y_1,\ldots,y_v)\sim(x_1,\ldots,x_v)} \Pr[A_{(y_1,\ldots,y_v)} | A_{(x_1,\ldots,x_v)}]$$

There are $v!/a = O(1)$ terms with $\{y_1,\ldots,y_v\} = \{x_1,\ldots,x_v\}$ and for each the conditional probability is at most one (actually, at most p), thus contributing $O(1) = o(E[X])$ to Δ^*. When $\{y_1,\ldots,y_v\} \cap \{x_1,\ldots,x_v\}$ has i elements, $2 \le i \le v-1$ the argument of Theorem 4.2 gives that the contribution to Δ^* is $o(E[X])$. Altogether $\Delta^* = o(E[X])$ and we apply Corollary 3.5 $\square$

general, complicated due to the overlapping of potential copies of H.) Let X_S be the indicator random variable for A_S and

$$X = \sum_{|S|=v} X_S$$

so that A holds if and only if $X > 0$. Linearity of Expectation gives

$$E[X] = \sum_{|S|=v} E[X_S] = \binom{n}{v} \Pr[A_S] = \Theta(n^v p^e)$$

If $p \ll n^{-v/e}$ then $E[X] = o(1)$ so $X = 0$ almost always.

Now assume $p \gg n^{-v/e}$ so that $E[X] \to \infty$ and consider the Δ^* of Corollary 3.5 (All v-sets look the same so the X_S are symmetric.) Here $S \sim T$ if and only if $S \neq T$ and S, T have common edges - i.e., if and only if $|S \cap T| = i$ with $2 \leq i \leq v-1$. Let S be fixed. We split

$$\Delta^* = \sum_{T \sim S} \Pr[A_T|A_S] = \sum_{i=2}^{v-1} \sum_{|T \cap S|=i} \Pr[A_T|A_S]$$

For each i there are $O(n^{v-i})$ choices of T. Fix S, T and consider $\Pr[A_T|A_S]$. There are $O(1)$ possible copies of H on T. Each has - since, critically, H is balanced - at most $\frac{ie}{v}$ edges with both vertices in S and thus at least $e - \frac{ie}{v}$ other edges. Hence

$$\Pr[A_T|A_S] = O(p^{e-\frac{ie}{v}})$$

and

$$\begin{aligned} \Delta^* &= \textstyle\sum_{i=2}^{v-1} O(n^{v-i} p^{e-\frac{ie}{v}}) \\ &= \textstyle\sum_{i=2}^{v-1} O((n^v p^e)^{1-\frac{i}{v}}) \\ &= \textstyle\sum_{i=2}^{v-1} o(n^v p^e) \\ &= o(E[X]) \end{aligned}$$

since $n^v p^e \to \infty$. Hence Corollary 3.5 applies. □

Theorem 4.3 In the notation of Theorem 4.2 if H is *not* balanced then $p = n^{-v/e}$ is *not* the threshold function for A.

Proof. Let H_1 be a subgraph of H with v_1 vertices, e_1 edges and $e_1/v_1 > e/v$. Let α satisfy $v/e < \alpha < v_1/e_1$ and set $p = n^{-\alpha}$. The expected number of copies of H_1 is then $o(1)$ so almost always $G(n,p)$ contains no copy of H_1. But if it contains no copy of H_1 then it surely can contain no copy of H. □

The threshold function for the property of containing a copy of H, for general H, was examined in the original papers of Erdős and Rényi. Let H_1

$\rho(H)$.
Examples. K_4 and, in general, K_k are strictly balanced. The graph

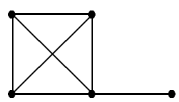

is not balanced as it has density 7/5 while the subgraph K_4 has density 3/2. The graph

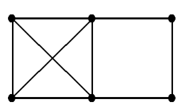

is balanced but not strictly balanced as it and its subgraph K_4 have density 3/2.
Theorem 4.2 Let H be a balanced graph with v vertices and e edges. Let $A(G)$ be the event that H is a subgraph (not necessarily induced) of G. Then $p = n^{-v/e}$ is the threshold function for A.
Proof. We follow the argument of Theorem 4.1 For each v-set S let A_S be the event that $G|_S$ contains H as a subgraph. Then

$$p^e \leq \Pr[A_S] \leq v!p^e$$

(Any particular placement of H has probability p^e of occuring and there are at most $v!$ possible placements. The precise calculation of $\Pr[A_S]$ is, in

4 Appearance of Small Subgraphs

What is the threshold function for the appearance of a given graph H. This problem was solved in the original papers of Erdős and Rényi. We begin with an instructive special case.

Theorem 4.1 The property $\omega(G) \geq 4$ has threshold function $n^{-2/3}$.

Proof. For every 4-set S of vertices in $G(n,p)$ let A_S be the event "S is a clique" and X_S its indicator random variable. Then

$$E[X_S] = \Pr[A_S] = p^6$$

as six different edges must all lie in $G(n,p)$. Set

$$X = \sum_{|S|=4} X_S$$

so that X is the number of 4-cliques in G and $\omega(G) \geq 4$ if and only if $X > 0$. Linearity of Expectation gives

$$E[X] = \sum_{|S|=4} E[X_S] = \binom{n}{4} p^6 \sim \frac{n^4 p^6}{24}$$

When $p(n) \ll n^{-2/3}$, $E[X] = o(1)$ and so $X = 0$ almost surely.

Now suppose $p(n) \gg n^{-2/3}$ so that $E[X] \to \infty$ and consider the Δ^* of Corollary 3.5. (All 4-sets "look the same" so that the X_S are symmetric.) Here $S \sim T$ if and only if $S \neq T$ and S, T have common edges - i.e., if and only if $|S \cap T| = 2$ or 3. Fix S. There are $O(n^2)$ sets T with $|S \cap T| = 2$ and for each of these $\Pr[A_T|A_S] = p^5$. There are $O(n)$ sets T with $|S \cap T| = 3$ and for each of these $\Pr[A_T|A_S] = p^3$. Thus

$$\Delta^* = O(n^2p^5) + O(np^3) = o(n^4p^6) = o(E[X])$$

since $p \gg n^{-2/3}$. Corollary 3.5 therefore applies and $X > 0$, i.e., there *does* exist a clique of size 4, almost always. □

The proof of Theorem 4.1 appears to require a fortuitous calculation of Δ^*. The following definitions will allow for a description of when these calculations work out.

Definitions. Let H be a graph with v vertices and e edges. We call $\rho(H) = e/v$ the *density* of H. We call H *balanced* if every subgraph H' has $\rho(H') \leq \rho(H)$. We call H *strictly balanced* if every proper subgraph H' has $\rho(H') <$

and thus in asymptotic terms we actually have the following stronger assertion:
Corollary 3.3

$$\text{If } Var[X] = o(E[X]^2) \text{ then } X \sim E[X] \text{ a.a.}$$

Suppose again $X = X_1 + \ldots + X_m$ where X_i is the indicator random variable for event A_i. For indices i, j write $i \sim j$ if $i \neq j$ and the events A_i, A_j are not independent. We set (the sum over ordered pairs)

$$\Delta = \sum_{i \sim j} \Pr[A_i \wedge A_j]$$

Note that when $i \sim j$

$$Cov[X_i, X_j] = E[X_i X_j] - E[X_i]E[X_j] \leq E[X_i X_j] = \Pr[A_i \wedge A_j]$$

and that when $i \neq j$ and not $i \sim j$ then $Cov[X_i, X_j] = 0$. Thus

$$Var[X] \leq E[X] + \Delta$$

Corollary 3.4. If $E[X] \to \infty$ and $\Delta = o(E[X]^2)$ then $X > 0$ almost always. Furthermore $X \sim E[X]$ almost always.

Let us say $X_1, \ldots, X_m$ are *symmetric* if for every $i \neq j$ there is an automorphism of the underlying probability space that sends event A_i to event A_j. Examples will appear in the next section. In this instance we write

$$\Delta = \sum_{i \sim j} \Pr[A_i \wedge A_j] = \sum_i \Pr[A_i] \sum_{j \sim i} \Pr[A_j | A_i]$$

and note that the inner summation is independent of i. We set

$$\Delta^* = \sum_{j \sim i} \Pr[A_j | A_i]$$

where i is any fixed index. Then

$$\Delta = \sum_i \Pr[A_i] \Delta^* = \Delta^* \sum_i \Pr[A_i] = \Delta^* E[X]$$

Corollary 3.5. If $E[X] \to \infty$ and $\Delta^* = o(E[X])$ then $X > 0$ almost always. Furthermore $X \sim E[X]$ almost always.

The condition of Corollary 3.5 has the intuitive sense that conditioning on any specific A_i holding does not substantially increase the expected number $E[X]$ of events holding.

The reality is that the B_S are not mutually independent though when $|S \cap T| \leq 1$, B_S and B_T are mutually independent. This is quite a typical situation in the study of random graphs in which we must deal with events that are "almost", but not precisely, mutual independent.

3 Variance

Here we introduce the Variance in a form that is particularly suited to the study of random graphs. The expressions Δ and Δ^* defined in this section will appear often in these notes.

Let X be a nonnegative integral valued random variable and suppose we want to bound $\Pr[X=0]$ given the value $\mu = E[X]$. If $\mu < 1$ we may use the inequality

$$\Pr[X > 0] \leq E[X]$$

so that if $E[X] \rightarrow 0$ then $X = 0$ almost always. (Here we are imagining an infinite sequence of X dependent on some parameter n going to infinity.) But now suppose $E[X] \rightarrow \infty$. It does *not* necessarily follow that $X > 0$ almost always. For example, let X be the number of deaths due to nuclear war in the twelve months after reading this paragraph. Calculation of $E[X]$ can make for lively debate but few would deny that it is quite large. Yet we may believe - or hope - that $Pr[X \neq 0]$ is very close to zero. We can sometimes deduce $X > 0$ almost always if we have further information about $Var[X]$.

Theorem 3.1

$$\Pr[X = 0] \leq \frac{Var[X]}{E[X]^2}$$

Proof. Set $\lambda = \mu/\sigma$ in Chebyschev's Inequality. Then

$$\Pr[X = 0] \leq \Pr[|X - \mu| \geq \lambda\sigma] \leq \frac{1}{\lambda^2} = \frac{\sigma^2}{\mu^2} \; \Box$$

We generally apply this result in asymptotic terms.

Corollary 3.2

$$\text{If } Var[X] = o(E[X]^2) \text{ then } X > 0 \text{ a.a.}$$

The proof of the Theorem actually gives that for any $\epsilon > 0$

$$\Pr[|X - E[X]| \geq \epsilon E[X]] \leq \frac{Var[X]}{\epsilon^2 E[X]^2}$$

This suggests the parametrization $p = c/n$. Then

$$\lim_{n\to\infty} E[X] = \lim_{n\to\infty} \binom{n}{3} p^3 = c^3/6$$

We shall see that the distribution of X is asymptotically Poisson. In particular

$$\lim_{n\to\infty} \Pr[G(n,p) \models A] = \lim_{n\to\infty} \Pr[X = 0] = e^{-c^3/6}$$

Note that

$$\lim_{c\to 0} e^{-c^3/6} = 1$$

$$\lim_{c\to\infty} e^{-c^3/6} = 0$$

When $p = 10^{-6}/n$, $G(n,p)$ is very unlikely to have triangles and when $p = 10^6/n$, $G(n,p)$ is very likely to have triangles. In the dynamic view the first triangles almost always appear at $p = \Theta(1/n)$. If we take a function such as $p(n) = n^{-.9}$ with $p(n) \gg n^{-1}$ then $G(n,p)$ will almost always have triangles. Occasionally we will abuse notation and say, for example, that $G(n, n^{-.9})$ contains a triangle - this meaning that the probability that it contains a triangle approaches 1 as n approaches infinity. Similarly, when $p(n) \ll n^{-1}$, for example, $p(n) = 1/(n \ln n)$, then $G(n,p)$ will almost always not contain a triangle and we abuse notation and say that $G(n, 1/(n \ln n))$ is trianglefree. It was a central observation of Erdős and Rényi that many natural graph theoretic properties become true in a very narrow range of p. They made the following key definition.
Definition. $r(n)$ is called a *threshold function* for a graph theoretic property A if
(i) When $p(n) \ll r(n), \lim_{n\to\infty} \Pr[G(n,p) \models A] = 0$
(ii) When $p(n) \gg r(n), \lim_{n\to\infty} \Pr[G(n,p) \models A] = 1$
or visa versa.

In our example, $1/n$ is a threshold function for A. Note that the threshold function, when one exists, is not unique. We could equally have said that $10/n$ is a threshold function for A.

Lets approach the problem of $G(n, c/n)$ being trianglefree once more. For every set S of three vertices let B_S be the event that S is a triangle. Then $\Pr[B_S] = p^3$. Then "trianglefreeness" is precisely the conjunction $\wedge \overline{B}_S$ over all S. If the B_S were mutually independent then we *would* have

$$\Pr[\wedge \overline{B}_S] = \prod [\overline{B}_S] = (1-p^3)^{\binom{n}{3}} \sim e^{-\binom{n}{3}p^3} \to e^{-c^3/6}$$

Lecture 1: Basics

1 What is a Random Graph

Let n be a positive integer, $0 \le p \le 1$. The random graph $G(n,p)$ is a probability space over the set of graphs on the vertex set $\{1,\ldots,n\}$ determined by

$$\Pr[\{i,j\} \in G] = p$$

with these events mutually independent.

Random Graphs is an active area of research which combines probability theory and graph theory. The subject began in 1960 with the monumental paper *On the Evolution of Random Graphs* by Paul Erdős and Alfred Rényi. The book *Random Graphs* by Béla Bollobás is the standard source for the field.

There is a compelling dynamic model for random graphs. For all pairs i,j let $x_{i,j}$ be selected uniformly from $[0,1]$, the choices mutually independent. Imagine p going from 0 to 1. Originally, all potential edges are "off". The edge from i to j (which we may imagine as a neon light) is turned on when p reaches $x_{i,j}$ and then stays on. At $p = 1$ all edges are "on". At time p the graph of all "on" edges has distribution $G(n,p)$. As p increases $G(n,p)$ *evolves* from empty to full.

In their original paper Erdős and Rényi let $G(n,e)$ be the random graph with n vertices and precisely e edges. Again there is a dynamic model: Begin with no edges and add edges randomly one by one until the graph becomes full. Generally $G(n,e)$ will have very similar properties as $G(n,p)$ with $p \sim \frac{e}{\binom{n}{2}}$. We will work on the probability model exclusively.

2 Threshold Functions

The term "the random graph" is, strictly speaking, a misnomer. $G(n,p)$ is a probability space over graphs. Given any graph theoretic property A there will be a probability that $G(n,p)$ satisfies A, which we write $\Pr[G(n,p) \models A]$. When A is monotone $\Pr[G(n,p) \models A]$ is a monotone function of p. As an instructive example, let A be the event "G is triangle free". Let X be the number of triangles contained in $G(n,p)$. Linearity of expectation gives

$$E[X] = \binom{n}{3} p^3$$