11 Random Graphs

At various points in this book, we already encountered the following fundamental theorem of Erdős: *for every integer k there is a graph G with $g(G) > k$ and $\chi(G) > k$.* In plain English: there exist graphs combining arbitrarily large girth with arbitrarily high chromatic number.

How could one prove such a theorem? The standard approach would be to construct a graph with those two properties, possibly in steps by induction on k. However, this is anything but straightforward: the global nature of the second property forced by the first, namely, that the graph should have high chromatic number 'overall' but be acyclic (and hence 2-colourable) locally, flies in the face of any attempt to build it up, constructively, from smaller pieces that have the same or similar properties.

In his pioneering paper of 1959, Erdős took a radically different approach: for each n he defined a probability space on the set of graphs with n vertices, and showed that, for some carefully chosen probability measures, the probability that an n-vertex graph has both of the above properties is positive for all large enough n.

This approach, now called the *probabilistic method*, has since unfolded into a sophisticated and versatile proof technique, in graph theory as much as in other branches of discrete mathematics. The theory of *random graphs* is now a subject in its own right. The aim of this chapter is to offer an elementary but rigorous introduction to random graphs: no more than is necessary to understand its basic concepts, ideas and techniques, but enough to give an inkling of the power and elegance hidden behind the calculations.

Erdős's theorem asserts the existence of a graph with certain properties: it is a perfectly ordinary assertion showing no trace of the randomness employed in its proof. There are also results in random graphs that are generically random even in their statement: these are theorems about *almost all* graphs, a notion we shall meet in Section 11.3. In the last section, we give a detailed proof of a theorem of Erdős and Rényi that illustrates a proof technique frequently used in random graphs, the so-called *second moment method*.

11.1 The notion of a random graph

V $\mathcal{G}$

Let V be a fixed set of n elements, say $V = \{0, \dots, n-1\}$. Our aim is to turn the set $\mathcal{G}$ of all graphs on V into a probability space, and then to consider the kind of questions typically asked about random objects: What is the probability that a graph $G \in \mathcal{G}$ has this or that property? What is the expected value of a given invariant on G, say its expected girth or chromatic number?

p q

Intuitively, we should be able to generate G randomly as follows. For each $e \in [V]^2$ we decide by some random experiment whether or not e shall be an edge of G; these experiments are performed independently, and for each the probability of success – i.e. of accepting e as an edge for G – is equal to some fixed[1] number $p \in [0,1]$. Then if G_0 is some fixed graph on V, with m edges say, the elementary event $\{G_0\}$ has a probability of $p^m q^{\binom{n}{2}-m}$ (where $q := 1-p$): with this probability, our randomly generated graph G is this particular graph G_0. (The probability that G is *isomorphic* to G_0 will usually be greater.) But if the probabilities of all the elementary events are thus determined, then so is the entire probability measure of our desired space $\mathcal{G}$. Hence all that remains to be checked is that such a probability measure on $\mathcal{G}$, one for which all individual edges occur independently with probability p, does indeed exist.[2]

Ω_e $\mathbb{P}_e$ $\mathcal{G}(n,p)$

In order to construct such a measure on $\mathcal{G}$ formally, we start by defining for every potential edge $e \in [V]^2$ its own little probability space $\Omega_e := \{0_e, 1_e\}$, choosing $\mathbb{P}_e(\{1_e\}) := p$ and $\mathbb{P}_e(\{0_e\}) := q$ as the probabilities of its two elementary events. As our desired probability space $\mathcal{G} = \mathcal{G}(n,p)$ we then take the product space

Ω

$$\Omega := \prod_{e \in [V]^2} \Omega_e .$$

[1] Often, the value of p will depend on the cardinality n of the set V on which our random graphs are generated; thus, p will be the value $p = p(n)$ of some function $n \mapsto p(n)$. Note, however, that V (and hence n) is fixed for the definition of $\mathcal{G}$: for each n separately, we are constructing a probability space of the graphs G on $V = \{0, \dots, n-1\}$, and within each space the probability that $e \in [V]^2$ is an edge of G has the same value for all e.

[2] Any reader ready to believe this may skip ahead now to the end of Proposition 11.1.1, without missing anything.

Thus, formally, an element of Ω is a map ω assigning to every $e \in [V]^2$ either 0_e or 1_e, and the probability measure $\mathbb{P}$ on Ω is the product measure of all the measures $\mathbb{P}_e$. In practice, of course, we identify ω with the graph G on V whose edge set is $\mathbb{P}$

$$E(G) = \{\, e \mid \omega(e) = 1_e \,\}\,,$$

and call G a *random graph* on V with edge probability p. *random graph*

Following standard probabilistic terminology, we may now call any set of graphs on V an *event* in $\mathcal{G}(n,p)$. In particular, for every $e \in [V]^2$ the set *event*

$$A_e := \{\, \omega \mid \omega(e) = 1_e \,\}$$ A_e

of all graphs G on V with $e \in E(G)$ is an event: the event that e is an edge of G. For these events, we can now prove formally what had been our guiding intuition all along:

Proposition 11.1.1. *The events A_e are independent and occur with probability p.*

Proof. By definition,

$$A_e = \{1_e\} \times \prod_{e' \neq e} \Omega_{e'}\,.$$

Since $\mathbb{P}$ is the product measure of all the measures $\mathbb{P}_e$, this implies

$$\mathbb{P}(A_e) = p \cdot \prod_{e' \neq e} 1 = p\,.$$

Similarly, if $\{e_1, \dots, e_k\}$ is any subset of $[V]^2$, then

$$\begin{aligned}
\mathbb{P}\left(A_{e_1} \cap \ldots \cap A_{e_k}\right) &= \mathbb{P}\Big(\{1_{e_1}\} \times \ldots \times \{1_{e_k}\} \times \prod_{e \notin \{e_1,\dots,e_k\}} \Omega_e\Big) \\
&= p^k \\
&= \mathbb{P}(A_{e_1}) \cdots \mathbb{P}(A_{e_k})\,.
\end{aligned}$$

□

As noted before, $\mathbb{P}$ is determined uniquely by the value of p and our assumption that the events A_e are independent. In order to calculate probabilities in $\mathcal{G}(n,p)$, it therefore generally suffices to work with these two assumptions: our concrete model for $\mathcal{G}(n,p)$ has served its purpose and will not be needed again.

As a simple example of such a calculation, consider the event that G contains some fixed graph H on a subset of V as a subgraph; let $|H| =: k$ and $\|H\| =: \ell$. The probability of this event $H \subseteq G$ is the product of the probabilities A_e over all the edges $e \in H$, so $\mathbb{P}[\,H \subseteq G\,] = p^\ell$. In k ℓ

contrast, the probability that H is an *induced* subgraph of G is $p^\ell q^{\binom{k}{2}-\ell}$: now the edges missing from H are required to be missing from G too, and they do so independently with probability q.

The probability p_H that G has an induced subgraph *isomorphic* to H is usually more difficult to compute: since the possible instances of H on subsets of V overlap, the events that they occur in G are not independent. However, the sum (over all k-sets $U \subseteq V$) of the probabilities $\mathbb{P}[\,H \cong G[U]\,]$ is always an upper bound for p_H, since p_H is the measure of the union of all those events. For example, if $H = \overline{K^k}$, we have the following trivial upper bound on the probability that G contains an induced copy of H:

[11.2.1]
[11.3.4]

Lemma 11.1.2. *For all integers n, k with $n \geqslant k \geqslant 2$, the probability that $G \in \mathcal{G}(n,p)$ has a set of k independent vertices is at most*

$$\mathbb{P}[\,\alpha(G) \geqslant k\,] \leqslant \binom{n}{k} q^{\binom{k}{2}}.$$

Proof. The probability that a fixed k-set $U \subseteq V$ is independent in G is $q^{\binom{k}{2}}$. The assertion thus follows from the fact that there are only $\binom{n}{k}$ such sets U. □

Analogously, the probability that $G \in \mathcal{G}(n,p)$ contains a K^k is at most

$$\mathbb{P}[\,\omega(G) \geqslant k\,] \leqslant \binom{n}{k} p^{\binom{k}{2}}.$$

Now if k is fixed, and n is small enough that these bounds for the probabilities $\mathbb{P}[\,\alpha(G) \geqslant k\,]$ and $\mathbb{P}[\,\omega(G) \geqslant k\,]$ sum to less than 1, then $\mathcal{G}$ contains graphs that have neither property: graphs which contain neither a K^k nor a $\overline{K^k}$ induced. But then any such n is a lower bound for the Ramsey number of k!

As the following theorem shows, this lower bound is quite close to the upper bound of 2^{2k-3} implied by the proof of Theorem 9.1.1:

Theorem 11.1.3. (Erdős 1947)
For every integer $k \geqslant 3$, the Ramsey number of k satisfies

$$R(k) > 2^{k/2}.$$

Proof. For $k = 3$ we trivially have $R(3) \geqslant 3 > 2^{3/2}$, so let $k \geqslant 4$. We show that, for all $n \leqslant 2^{k/2}$ and $G \in \mathcal{G}(n, \frac{1}{2})$, the probabilities $\mathbb{P}[\,\alpha(G) \geqslant k\,]$ and $\mathbb{P}[\,\omega(G) \geqslant k\,]$ are both less than $\frac{1}{2}$.

Since $p = q = \frac{1}{2}$, Lemma 11.1.2 and the analogous assertion for $\omega(G)$ imply the following for all $n \leqslant 2^{k/2}$ (use that $k! > 2^k$ for $k \geqslant 4$):

$$\begin{aligned}\mathbb{P}\,[\,\alpha(G) \geqslant k\,],\ \mathbb{P}\,[\,\omega(G) \geqslant k\,] &\leqslant \binom{n}{k}\left(\tfrac{1}{2}\right)^{\binom{k}{2}} \\ &< \left(n^k/2^k\right) 2^{-\frac{1}{2}k(k-1)} \\ &\leqslant \left(2^{k^2/2}/2^k\right) 2^{-\frac{1}{2}k(k-1)} \\ &= 2^{-k/2} \\ &< \tfrac{1}{2}\,.\end{aligned}$$

□

In the context of random graphs, most of the familiar graph invariants (like average degree, connectivity, chromatic number, and so on) may be interpreted as a non-negative *random variable* on $\mathcal{G}(n,p)$, a function

random variable

$$X\colon \mathcal{G}(n,p) \to [0,\infty)\,.$$

The *mean* or *expected* value of X is the number

mean

expectation

$$\mathbb{E}(X) := \sum_{G \in \mathcal{G}(n,p)} \mathbb{P}(\{G\}) \cdot X(G)\,.$$

$\mathbb{E}$

If X takes integers as values, we can compute $\mathbb{E}(X)$ alternatively by summing over these values k:

$$\mathbb{E}(X) = \sum_{k \geqslant 1} \mathbb{P}\,[\,X \geqslant k\,] = \sum_{k \geqslant 1} k \cdot \mathbb{P}\,[\,X = k\,]\,.$$

Note also that the operator $\mathbb{E}$, the *expectation*, is linear: we have $\mathbb{E}(X+Y) = \mathbb{E}(X) + \mathbb{E}(Y)$ and $\mathbb{E}(\lambda X) = \lambda\mathbb{E}(X)$ for any two random variables X, Y on $\mathcal{G}(n,p)$ and $\lambda \geqslant 0$.

Since our probability spaces are finite, the expectation can often be computed by a simple application of *double counting*, a standard combinatorial technique we met before in the proof of Corollary 4.2.10. For example, if X is a random variable on $\mathcal{G}(n,p)$ that counts the number of subgraphs of G in some fixed set $\mathcal{H}$ of graphs on V, then $\mathbb{E}(X)$, by definition, counts the number of pairs (G, H) such that $H \in \mathcal{H}$ and $H \subseteq G$, each weighted with the probability $\mathbb{P}(\{G\})$. Algorithmically, we compute $\mathbb{E}(X)$ by going through the graphs $G \in \mathcal{G}(n,p)$ in an 'outer loop' and performing, for each G, an 'inner loop' that runs through the graphs $H \in \mathcal{H}$ and counts '$\mathbb{P}(\{G\})$' whenever $H \subseteq G$. Alternatively, we may count the same set of weighted pairs with H in the outer and G in the inner loop. This amounts to adding up, over all $H \in \mathcal{H}$, the probabilities $\mathbb{P}\,[\,H \subseteq G\,]$:

$$\mathbb{E}(X) = \sum_{G \in \mathcal{G}(n,p)} \big|\{\,H \in \mathcal{H} : H \subseteq G\,\}\big| \cdot \mathbb{P}(\{G\}) = \sum_{H \in \mathcal{H}} \mathbb{P}\,[\,H \subseteq G\,].$$

To illustrate this once in detail, and to introduce the probabilistic terminology commonly used at this point, let us compute the expected number of cycles of some given length $k \geqslant 3$ in a random graph $G \in \mathcal{G}(n,p)$. (We shall also need this for our proof of Erdős's theorem in Section 11.2.) Let $X : \mathcal{G}(n,p) \to \mathbb{N}$ be the random variable that assigns to a random graph G its number of k-cycles, the number of subgraphs isomorphic to C^k.

X

How many potential such cycles are there? In other words, how large is the set $\mathcal{C}_k$ of all k-cycles with vertices in V? Since there are

$\mathcal{C}_k$

$(n)_k$

$$(n)_k := n\,(n-1)(n-2)\cdots(n-k+1)$$

ways of choosing a sequence of k distinct vertices in V, and each k-cycle is identified by $2k$ of those sequences, clearly

$$|\mathcal{C}_k| = (n)_k/2k\,. \tag{1}$$

[11.2.2] [11.4.3]

Lemma 11.1.4. *The expected number of k-cycles in $G \in \mathcal{G}(n,p)$ is*

$$\mathbb{E}(X) = \frac{(n)_k}{2k}\,p^k.$$

Proof. Consider for every fixed $C \in \mathcal{C}_k$ its *indicator random variable* $X_C : \mathcal{G}(n,p) \to \{0,1\}$, defined by

$$X_C\colon\; G \mapsto \begin{cases} 1 & \text{if } C \subseteq G;\\ 0 & \text{otherwise.}\end{cases}$$

Since X_C takes only 1 as a positive value, its expectation $\mathbb{E}(X_C)$ equals the measure $\mathbb{P}[\,X_C = 1\,]$ of the set of all graphs in $\mathcal{G}(n,p)$ that contain C. But this is just the probability that $C \subseteq G$:

$$\mathbb{E}(X_C) = \mathbb{P}[\,C \subseteq G\,] = p^k. \tag{2}$$

Our random variable X assigns to every graph G its number of k-cycles. Hence

$$X(G) = \sum_{C \in \mathcal{C}_k} X_C(G)$$

for every G, or $X = \sum X_C$ for short. Since the expectation is linear, applying this with (1) and (2) yields

$$\mathbb{E}(X) = \sum_{C \in \mathcal{C}_k} \mathbb{E}(X_C) = \sum_{C \in \mathcal{C}_k} \mathbb{P}[\,C \subseteq G\,] = \frac{(n)_k}{2k}\,p^k$$

as claimed. □

Computing the mean of a random variable X can be a simple and effective way to establish the existence of a graph G such that $X(G) < a$ for some fixed $a > 0$ and, moreover, G has some desired property $\mathcal{P}$. Indeed, if the expected value of X is small, then $X(G)$ cannot be large for more than a few graphs in $\mathcal{G}(n,p)$, because $X(G) \geqslant 0$ for all $G \in \mathcal{G}(n,p)$. Hence X must be small for many graphs in $\mathcal{G}(n,p)$, and it is reasonable to expect that among these we may find one with the desired property $\mathcal{P}$.

This simple idea lies at the heart of countless non-constructive existence proofs using random graphs, including the proof of Erdős's theorem presented in the next section. Quantified, it takes the form of the following lemma, whose proof follows at once from the definition of the expectation and the additivity of $\mathbb{P}$:

[11.2.2]
[11.4.1]
[11.4.3]

Lemma 11.1.5. (Markov's Inequality)
Let $X \geqslant 0$ be a random variable on $\mathcal{G}(n,p)$ and $a > 0$. Then

$$\mathbb{P}\,[\,X \geqslant a\,] \leqslant \mathbb{E}(X)/a\,.$$

Proof.

$$\mathbb{E}(X) = \sum_{G \in \mathcal{G}(n,p)} \mathbb{P}(\{G\}) \cdot X(G) \;\geqslant \sum_{\substack{G \in \mathcal{G}(n,p) \\ X(G) \geqslant a}} \mathbb{P}(\{G\}) \cdot a = \mathbb{P}\,[\,X \geqslant a\,] \cdot a\,.$$

□

11.2 The probabilistic method

Very roughly, the *probabilistic method* in discrete mathematics has developed from the following idea. In order to prove the existence of an object with some desired property, one defines a probability space on some larger – and certainly non-empty – class of objects, and then shows that an element of this space has the desired property with positive probability. The 'objects' inhabiting this probability space may be of any kind: partitions or orderings of the vertices of some fixed graph arise as naturally as mappings, embeddings and, of course, graphs themselves. In this section, we illustrate the probabilistic method by giving a detailed account of one of its earliest results: of Erdős's classic theorem on large girth and chromatic number (Theorem 5.2.5).

Erdős's theorem says that, given any positive integer k, there is a graph G with girth $g(G) > k$ and chromatic number $\chi(G) > k$. Let us call cycles of length at most k *short*, and sets of $|G|/k$ or more vertices *short*
big. For a proof of Erdős's theorem, it suffices to find a graph G without *big/small*
short cycles and without big independent sets of vertices: then the colour classes in any vertex colouring of G are *small* (not big), so we need more than k colours to colour G.

How can we find such a graph G? If we choose p small enough, then a random graph in $\mathcal{G}(n,p)$ is unlikely to contain any (short) cycles. If we choose p large enough, then G is unlikely to have big independent vertex sets. So the question is: do these two ranges of p overlap, that is, can we choose p so that, for some n, it is both small enough to give $\mathbb{P}[\,g \leqslant k\,] < \frac{1}{2}$ and large enough for $\mathbb{P}[\,\alpha \geqslant n/k\,] < \frac{1}{2}$? If so, then $\mathcal{G}(n,p)$ will contain at least one graph without either short cycles or big independent sets.

Unfortunately, such a choice of p is impossible: the two ranges of p do not overlap! As we shall see in Section 11.4, we must keep p below n^{-1} to make the occurrence of short cycles in G unlikely – but for any such p there will most likely be no cycles in G at all (Exercise 18), so G will be bipartite and hence have at least $n/2$ independent vertices.

But all is not lost. In order to make big independent sets unlikely, we shall fix p above n^{-1}, at $n^{\epsilon-1}$ for some $\epsilon > 0$. Fortunately, though, if ϵ is small enough then this will produce only *few* short cycles in G, even compared with n (rather than, more typically, with n^k). If we then delete a vertex in each of those cycles, the graph H obtained will have no short cycles, and its independence number $\alpha(H)$ will be at most that of G. Since H is not much smaller than G, its chromatic number will thus still be large, so we have found a graph with both large girth and large chromatic number.

To prepare for the formal proof of Erdős's theorem, we first show that an edge probability of $p = n^{\epsilon-1}$ is indeed always large enough to ensure that $G \in \mathcal{G}(n,p)$ 'almost surely' has no big independent set of vertices. More precisely, we prove the following stronger assertion:

Lemma 11.2.1. *Let $k > 0$ be an integer, and let $p = p(n)$ be a function of n such that $p(n) \geqslant 16k^2/n$ for n large. Then*

$$\lim_{n\to\infty} \mathbb{P}[\,\alpha \geqslant \tfrac{1}{2}n/k\,] = 0\,.$$

(11.1.2) *Proof.* For all integers $n \geqslant r \geqslant 2$ and $G \in \mathcal{G}(n,p)$, Lemma 11.1.2 implies[3]

$$\mathbb{P}[\,\alpha \geqslant r\,] \leqslant \binom{n}{r} q^{\binom{r}{2}} \leqslant 2^n q^{\binom{r}{2}} \leqslant 2^n e^{-p\binom{r}{2}}.$$

Hence if $p \geqslant 16k^2/n$ and $r \geqslant \frac{1}{2}n/k \geqslant 2$, then

$$\mathbb{P}[\,\alpha \geqslant r\,] \leqslant 2^n e^{-pr^2/4} \leqslant 2^n e^{-pn^2/16k^2} \leqslant 2^n e^{-n} \underset{n\to\infty}{\longrightarrow} 0\,.$$

As α is an integer and thus $\mathbb{P}[\,\alpha \geqslant r\,] = \mathbb{P}[\,\alpha \geqslant \frac{1}{2}n/k\,]$ for $r = \lceil \frac{1}{2}n/k \rceil$, this implies the assertion. □

[3] To see the second inequality, count the subsets of an n-set. For the third, note that $1-p \leqslant e^{-p}$ for all p: compare the functions $x \mapsto e^x$ and $x \mapsto x+1$ for $x = -p$.

We are now ready to prove Theorem 5.2.5, which we restate:

Theorem 11.2.2. (Erdős 1959) [9.2.3]
For every integer k there exists a graph H with girth $g(H) > k$ and chromatic number $\chi(H) > k$.

(11.1.5)
Proof. Assume that $k \geqslant 3$, fix ϵ with $0 < \epsilon < 1/k$, and let $p := n^{\epsilon - 1}$. Let (11.1.4)
$X(G)$ denote the number of short cycles in a random graph $G \in \mathcal{G}(n,p)$, p, ϵ, X
i.e. its number of cycles of length at most k.

By Lemma 11.1.4, we have

$$\mathbb{E}(X) = \sum_{i=3}^{k} \frac{(n)_i}{2i} p^i \leqslant \tfrac{1}{2} \sum_{i=3}^{k} n^i p^i \leqslant \tfrac{1}{2}(k-2)\, n^k p^k \,;$$

note that $(np)^i \leqslant (np)^k$, because $np = n^\epsilon \geqslant 1$. By Lemma 11.1.5,

$$\begin{aligned} \mathbb{P}\,[\, X \geqslant n/2 \,] &\leqslant \mathbb{E}(X)/(n/2) \\ &\leqslant (k-2)\, n^{k-1} p^k \\ &= (k-2)\, n^{k-1} n^{(\epsilon-1)k} \\ &= (k-2)\, n^{k\epsilon - 1}. \end{aligned}$$

As $k\epsilon - 1 < 0$ by our choice of ϵ, this implies that

$$\lim_{n \to \infty} \mathbb{P}\,[\, X \geqslant n/2 \,] = 0\,.$$

Let n be large enough that $\mathbb{P}\,[\, X \geqslant n/2 \,] < \frac{1}{2}$ and $\mathbb{P}\,[\, \alpha \geqslant \frac{1}{2} n/k \,] < \frac{1}{2}$; n
the latter is possible by our choice of p and Lemma 11.2.1. Then there is a graph $G \in \mathcal{G}(n,p)$ with fewer than $n/2$ short cycles and $\alpha(G) < \frac{1}{2} n/k$. From each of those cycles delete a vertex, and let H be the graph obtained. Then $|H| \geqslant n/2$ and H has no short cycles, so $g(H) > k$. By definition of G,

$$\chi(H) \geqslant \frac{|H|}{\alpha(H)} \geqslant \frac{n/2}{\alpha(G)} > k\,.$$

□

Corollary 11.2.3. *There are graphs with arbitrarily large girth and arbitrarily large values of the invariants κ, ε and δ.*

Proof. Apply Lemma 5.2.3 and Theorem 1.4.3. □ (1.4.3)
(5.2.3)

11.3 Properties of almost all graphs

Recall that a *graph property* is a class of graphs that is closed under isomorphism, one that contains with every graph G also the graphs isomorphic to G. If $p = p(n)$ is a fixed function (possibly constant), and $\mathcal{P}$ is a graph property, we may ask how the probability $\mathbb{P}[\, G \in \mathcal{P}\,]$ behaves for $G \in \mathcal{G}(n,p)$ as $n \to \infty$. If this probability tends to 1, we say that $G \in \mathcal{P}$ for *almost all* (or *almost every*) $G \in \mathcal{G}(n,p)$, or that $G \in \mathcal{P}$ *almost surely*; if it tends to 0, we say that *almost no* $G \in \mathcal{G}(n,p)$ has the property $\mathcal{P}$. (For example, in Lemma 11.2.1 we proved that, for a certain p, almost no $G \in \mathcal{G}(n,p)$ has a set of more than $\frac{1}{2}n/k$ independent vertices.)

almost all etc.

To illustrate the new concept let us show that, for constant p, every fixed abstract[4] graph H is an induced subgraph of almost all graphs:

Proposition 11.3.1. *For every constant $p \in (0,1)$ and every graph H, almost every $G \in \mathcal{G}(n,p)$ contains an induced copy of H.*

Proof. Let H be given, and $k := |H|$. If $n \geqslant k$ and $U \subseteq \{0,\dots,n-1\}$ is a fixed set of k vertices of G, then $G[U]$ is isomorphic to H with a certain probability $r > 0$. This probability r depends on p, but not on n (why not?). Now G contains a collection of $\lfloor n/k \rfloor$ disjoint such sets U. The probability that none of the corresponding graphs $G[U]$ is isomorphic to H is $(1-r)^{\lfloor n/k \rfloor}$, since these events are independent by the disjointness of the edges sets $[U]^2$. Thus

$$\mathbb{P}[\, H \not\subseteq G \text{ induced}\,] \leqslant (1-r)^{\lfloor n/k \rfloor} \underset{n\to\infty}{\longrightarrow} 0\,,$$

which implies the assertion. □

The following lemma is a simple device enabling us to deduce that quite a number of natural graph properties (including that of Proposition 11.3.1) are shared by almost all graphs. Given $i,j \in \mathbb{N}$, let $\mathcal{P}_{i,j}$ denote the property that the graph considered contains, for any disjoint vertex sets U, W with $|U| \leqslant i$ and $|W| \leqslant j$, a vertex $v \notin U \cup W$ that is adjacent to all the vertices in U but to none in W.

$\mathcal{P}_{i,j}$

Lemma 11.3.2. *For every constant $p \in (0,1)$ and $i,j \in \mathbb{N}$, almost every graph $G \in \mathcal{G}(n,p)$ has the property $\mathcal{P}_{i,j}$.*

[4] The word ‘abstract’ is used to indicate that only the isomorphism type of H is known or relevant, not its actual vertex and edge sets. In our context, it indicates that the word ‘subgraph’ is used in the usual sense of ‘isomorphic to a subgraph’.

Proof. For fixed U, W and $v \in G - (U \cup W)$, the probability that v is adjacent to all the vertices in U but to none in W, is

$$p^{|U|} q^{|W|} \geqslant p^i q^j .$$

Hence, the probability that no suitable v exists for these U and W, is

$$(1 - p^{|U|} q^{|W|})^{n-|U|-|W|} \leqslant (1 - p^i q^j)^{n-i-j}$$

(for $n \geqslant i + j$), since the corresponding events are independent for different v. As there are no more than n^{i+j} pairs of such sets U, W in $V(G)$ (encode sets U of fewer than i points as non-injective maps $\{0, \dots, i-1\} \to \{0, \dots, n-1\}$, etc.), the probability that some such pair has no suitable v is at most

$$n^{i+j} (1 - p^i q^j)^{n-i-j},$$

which tends to zero as $n \to \infty$ since $1 - p^i q^j < 1$. □

Corollary 11.3.3. *For every constant $p \in (0,1)$ and $k \in \mathbb{N}$, almost every graph in $\mathcal{G}(n,p)$ is k-connected.*

Proof. By Lemma 11.3.2, it is enough to show that every graph in $\mathcal{P}_{2,k-1}$ is k-connected. But this is easy: any graph in $\mathcal{P}_{2,k-1}$ has order at least $k+2$, and if W is a set of fewer than k vertices, then by definition of $\mathcal{P}_{2,k-1}$ any other two vertices x, y have a common neighbour $v \notin W$; in particular, W does not separate x from y. □

In the proof of Corollary 11.3.3, we showed substantially more than was asked for: rather than finding, for any two vertices $x, y \notin W$, some x–y path avoiding W, we showed that x and y have a common neighbour outside W; thus, all the paths needed to establish the desired connectivity could in fact be chosen of length 2. What seemed like a clever trick in this particular proof is in fact indicative of a more fundamental phenomenon for constant edge probabilities: by an easy result in logic, any statement about graphs expressed by quantifying over vertices only (rather than over sets or sequences of vertices)[5] is either almost surely true or almost surely false. All such statements, or their negations, are in fact immediate consequences of an assertion that the graph has property $\mathcal{P}_{i,j}$, for some suitable i, j.

As a last example of an 'almost all' result we now show that almost every graph has a surprisingly high chromatic number:

[5] In the terminology of logic: any first order sentence in the language of graph theory

Proposition 11.3.4. *For every constant $p \in (0,1)$ and every $\epsilon > 0$, almost every graph $G \in \mathcal{G}(n,p)$ has chromatic number*

$$\chi(G) > \frac{\log(1/q)}{2+\epsilon} \cdot \frac{n}{\log n}\,.$$

(11.1.2) *Proof.* For any fixed $n \geqslant k \geqslant 2$, Lemma 11.1.2 implies

$$\begin{aligned} \mathbb{P}\,[\,\alpha \geqslant k\,] &\leqslant \binom{n}{k}\, q^{\binom{k}{2}} \\ &\leqslant n^k q^{\binom{k}{2}} \\ &= q^{k \frac{\log n}{\log q} + \frac{1}{2}k(k-1)} \\ &= q^{\frac{k}{2}\left(-\frac{2\log n}{\log(1/q)} + k - 1\right)}. \end{aligned}$$

For

$$k := (2+\epsilon)\,\frac{\log n}{\log(1/q)}$$

the exponent of this expression tends to infinity with n, so the expression itself tends to zero. Hence, almost every $G \in \mathcal{G}(n,p)$ is such that in any vertex colouring of G no k vertices can have the same colour, so every colouring uses more than

$$\frac{n}{k} = \frac{\log(1/q)}{2+\epsilon} \cdot \frac{n}{\log n}$$

colours. □

By a result of Bollobás (1988), Proposition 11.3.4 is sharp in the following sense: if we replace ϵ by $-\epsilon$, then the lower bound given for χ turns into an upper bound.

We finish this section with a little gem, the one and only theorem about infinite random graphs. Let $\mathcal{G}(\aleph_0, p)$ be defined exactly like $\mathcal{G}(n,p)$ for $n = \aleph_0$, as the (product) space of random graphs on $\mathbb{N}$ whose edges are chosen independently with probability p.

As we saw in Lemma 11.3.2, the properties $\mathcal{P}_{i,j}$ hold almost surely for finite random graphs with constant edge probability. It will therefore hardly come as a surprise that an infinite random graph almost surely (which now has the usual meaning of 'with probability 1') has all these properties at once. However, in Chapter 8.3 we saw that, up to isomorphism, there is exactly one countable graph, the Rado graph R, that has property $\mathcal{P}_{i,j}$ for all $i,j \in \mathbb{N}$ simultaneously; this joint property was denoted as $(*)$ there. Combining these facts, we get the following rather bizarre result:

Theorem 11.3.5. (Erdős & Rényi 1963)
With probability 1, a random graph $G \in \mathcal{G}(\aleph_0, p)$ with $0 < p < 1$ is isomorphic to the Rado graph R.

Proof. Given fixed disjoint finite sets $U, W \subseteq \mathbb{N}$, the probability that a (8.3.1) vertex $v \notin U \cup W$ is not joined to $U \cup W$ as expressed in property $(*)$ of Chapter 8.3 (i.e., is not joined to all of U or is joined to some vertex in W) is some number $r < 1$ depending only on U and W. The probability that none of k given vertices v is joined to $U \cup W$ as in $(*)$ is r^k, which tends to 0 as $k \to \infty$. Hence the probability that all the (infinitely many) vertices outside $U \cup W$ fail to witness $(*)$ for these sets U and W is 0.

Now there are only countably many choices for U and W as above. Since the union of countably many sets of measure 0 again has measure 0, the probability that $(*)$ fails for *any* sets U and W is still 0. Therefore G satisfies $(*)$ with probability 1. By Theorem 8.3.1 this means that, almost surely, $G \cong R$. □

How can we make sense of the paradox that the result of infinitely many independent choices can be so predictable? The answer, of course, lies in the fact that the uniqueness of R holds only up to isomorphism. Now, constructing an automorphism for an infinite graph with property $(*)$ is a much easier task than finding one for a finite random graph, so in this sense the uniqueness is no longer that surprising. Viewed in this way, Theorem 11.3.5 expresses not a lack of variety in infinite random graphs but rather the abundance of symmetry that glosses over this variety when the graphs $G \in \mathcal{G}(\aleph_0, p)$ are viewed only up to isomorphism.

11.4 Threshold functions and second moments

The results we saw in Section 11.3 have an interesting common feature: the values of p played no role as long as they were constant, that is, independent of n. For example, if almost every graph in $\mathcal{G}(n, p)$ with $p = 0.99$ had the property considered, then the same was true for $p = 0.01$. How could this happen?

Such insensitivity of our random model to changes of p was certainly not intended. For most properties, however, the critical order of magnitude of p around which the property will 'just' occur or not occur simply lies below any constant value of p: it is a function of n tending to zero as $n \to \infty$. In the proof of Erdős's theorem, for example, this critical probability for the two properties we were trying to relate was $p(n) = 1/n$.

threshold function

Let us call a positive real function $t = t(n)$ a *threshold function* for a graph property $\mathcal{P}$ if the following holds for all $p = p(n)$ and $G \in \mathcal{G}(n,p)$:

$$\lim_{n\to\infty} \mathbb{P}\,[\,G \in \mathcal{P}\,] = \begin{cases} 0 & \text{if } p/t \to 0 \text{ as } n \to \infty \\ 1 & \text{if } p/t \to \infty \text{ as } n \to \infty. \end{cases}$$

If $\mathcal{P}$ has a threshold function t, then clearly any positive multiple ct of t is also a threshold function for $\mathcal{P}$; thus, threshold functions in the above sense are only ever unique up to a multiplicative constant.[6]

Bollobás & Thomason (1987) have shown that, trivial exceptions aside, all *increasing* graph properties have threshold functions, properties that are closed under the addition of edges. Properties of the form $\{\,G \mid G \supseteq H\,\}$, with H fixed, are a common example; we shall compute their threshold functions in this section.

For the purpose of computing its threshold function, it is convenient to cast the graph property $\mathcal{P}$ considered in the form

$$\mathcal{P} = \{\,G \mid X(G) \geqslant 1\,\}\,,$$

$X \geqslant 0$

where $X \geqslant 0$ is a suitable random variable on $\mathcal{G}(n,p)$. For example, we could take the indicator random variable of $\mathcal{P}$ on $\mathcal{G}(n,p)$. But other choices of X are allowed too; if $\mathcal{P}$ is connectedness, for example, we might have $X(G)$ count the number of spanning trees of G.

How could we prove that some given t is a threshold function of $\mathcal{P}$? Any such proof will consist of two parts: a proof that almost no $G \in \mathcal{G}(n,p)$ has $\mathcal{P}$ when p is small compared with t, and one showing that almost every G has $\mathcal{P}$ when p is large.

Since $X \geqslant 0$, we may use Markov's inequality for the first part of the proof and find an upper bound for $\mathbb{E}(X)$ instead of $\mathbb{P}\,[\,X \geqslant 1\,]$: if $\mathbb{E}(X)$ is much smaller than 1 then $X(G)$ can be at least 1 only for few $G \in \mathcal{G}(n,p)$, and for almost no G if $\mathbb{E}(X) \to 0$ as $n \to \infty$. Besides, the expectation is much easier to calculate than probabilities: without worrying about such things as independence or incompatibility of events, we may compute the expectation of a sum of random variables – for example, of indicator random variables – simply by adding up their individual expected values.

For the second part of the proof, things are more complicated. In order to show that $\mathbb{P}\,[\,X \geqslant 1\,]$ is large, it is not enough to bound $\mathbb{E}(X)$ from below: since X is not bounded above, $\mathbb{E}(X)$ may be large simply because X is very large on just a few graphs G – so X may still be zero

[6] Our notion of threshold reflects only the crudest interesting level of screening: for some properties, such as connectedness, one can define sharper thresholds where the constant factor is crucial.

for most $G \in \mathcal{G}(n,p)$.[7] In order to prove that $\mathbb{P}[X \geqslant 1] \to 1$, we thus have to show that this cannot happen, i.e., that X does not deviate a lot from its mean too often.

The following tool from probability theory achieves just that. As is customary, we write

$$\mu := \mathbb{E}(X)$$ μ

and define $\sigma \geqslant 0$ by setting

$$\sigma^2 := \mathbb{E}\big((X-\mu)^2\big)\,.$$ σ^2

This quantity σ^2 is called the *variance* or *second moment* of X; by definition, it is a measure of how much X deviates from its mean. Since $\mathbb{E}$ is linear, the defining term for σ^2 expands to

$$\sigma^2 = \mathbb{E}(X^2 - 2\mu X + \mu^2) = \mathbb{E}(X^2) - \mu^2.$$

Note that μ and σ^2 always refer to a random variable on some fixed probability space. In our setting, where we consider the spaces $\mathcal{G}(n,p)$, both quantities are functions of n.

The following lemma says exactly what we need: that X cannot deviate a lot from its mean too often.

Lemma 11.4.1. (Chebyshev's Inequality)
For all real $\lambda > 0$,

$$\mathbb{P}\big[\,|X-\mu| \geqslant \lambda\,\big] \;\leqslant\; \sigma^2/\lambda^2.$$

Proof. By Lemma 11.1.5 and definition of σ^2, (11.1.5)

$$\mathbb{P}\big[\,|X-\mu| \geqslant \lambda\,\big] = \mathbb{P}\big[\,(X-\mu)^2 \geqslant \lambda^2\,\big] \leqslant \sigma^2/\lambda^2.$$

□

For a proof that $X(G) \geqslant 1$ for almost all $G \in \mathcal{G}(n,p)$, Chebyshev's inequality can be used as follows:

Lemma 11.4.2. *If $\mu > 0$ for all large enough n, and $\sigma^2/\mu^2 \to 0$ as $n \to \infty$, then $X(G) > 0$ for almost all $G \in \mathcal{G}(n,p)$.*

---

[7] For some p between n^{-1} and $(\log n)n^{-1}$, for example, almost every $G \in \mathcal{G}(n,p)$ has an isolated vertex (and hence no spanning tree), but its expected number of spanning trees tends to infinity with n. See the Exercise 11 for details.

Proof. Any graph G with $X(G) = 0$ satisfies $|X(G) - \mu| = \mu$. Hence Lemma 11.4.1 implies with $\lambda := \mu$ that

$$\mathbb{P}\left[\, X = 0 \,\right] \;\leqslant\; \mathbb{P}\left[\, |X - \mu| \geqslant \mu \,\right] \;\leqslant\; \sigma^2/\mu^2 \underset{n\to\infty}{\longrightarrow} 0\,.$$

Since $X \geqslant 0$, this means that $X > 0$ almost surely, i.e. that $X(G) > 0$ for almost all $G \in \mathcal{G}(n,p)$. □

As the main result of this section, we now prove a theorem that will $\mathcal{P}_H$ give us a threshold function for all graph properties of the form $\mathcal{P}_H$, the property of containing a copy of a fixed graph H as a subgraph.

H; k, ℓ Let H be given, put $k := |H|$ and $\ell := \|H\|$, and assume that $\ell \geqslant 1$. Write $X(G)$ for the number of subgraphs of a graph G that are X isomorphic to H.

Given $n \in \mathbb{N}$, let $\mathcal{H}$ denote the set of all copies of H on subsets of $\{0, \dots, n-1\}$, the vertex set of the graphs in $\mathcal{G}(n,p)$:

$\mathcal{H}$
$$\mathcal{H} := \left\{ H' \mid H' \cong H,\; V(H') \subseteq \{0, \dots, n-1\} \right\}.$$

Given $H' \in \mathcal{H}$ and $G \in \mathcal{G}(n,p)$, we shall write $H' \subseteq G$ to express that H' itself – not just an isomorphic copy of H' – is a subgraph of G. As in the proof of Lemma 11.1.4, double counting gives

$$\mathbb{E}(X) = \sum_{H' \in \mathcal{H}} \mathbb{P}\left[\, H' \subseteq G \,\right]. \tag{1}$$

And for every fixed $H' \in \mathcal{H}$ we have

$$\mathbb{P}\left[\, H' \subseteq G \,\right] \;=\; p^{\ell}, \tag{2}$$

because $\|H\| = \ell$.

h Let h denote the number of isomorphic copies of H on a fixed k-set; clearly, $h \leqslant k!$. As there are $\binom{n}{k}$ possible vertex sets for the graphs in $\mathcal{H}$, we thus have

$$|\mathcal{H}| = \binom{n}{k} h \leqslant \binom{n}{k} k! \leqslant n^k. \tag{3}$$

Given a probability $p = p(n)$ and a candidate $t = t(n)$ for a threshold γ function, we write $\gamma := p/t$. Our first lemma treats the case of $\gamma \to 0$:

Lemma 11.4.3. *If $t = n^{-1/\varepsilon(H)}$, and p is such that $\gamma \to 0$ as $n \to \infty$, then almost no $G \in \mathcal{G}(n,p)$ lies in $\mathcal{P}_H$.*

Proof. Our aim is to find an upper bound for $\mathbb{E}(X)$, and to show that this tends to zero as $n \to \infty$. By our choice of t and definition of γ we have (11.1.5) (11.1.4)

$$p = \gamma t = \gamma n^{-k/\ell},$$

so

$$\mathbb{E}(X) \underset{(1,2)}{=} |\mathcal{H}|\, p^{\ell} \underset{(3)}{\leqslant} n^{k}(\gamma n^{-k/\ell})^{\ell} = \gamma^{\ell}.$$

Thus if $\gamma \to 0$ as $n \to \infty$, then

$$\mathbb{P}[\, G \in \mathcal{P}_H \,] \;=\; \mathbb{P}[\, X \geqslant 1 \,] \;\leqslant\; \mathbb{E}(X) \;\leqslant\; \gamma^{\ell} \underset{n\to\infty}{\longrightarrow} 0$$

by Markov's inequality (11.1.5). □

Unlike the function t in Lemma 11.4.3, our threshold function for $\mathcal{P}_H$ will not be expressed in terms of $\varepsilon(H)$, but of

$$\varepsilon'(H) := \max\{\, \varepsilon(H') \mid H' \subseteq H \,\}.$$ $\varepsilon'(H)$

Our second lemma treats the case of $\gamma \to \infty$:

Lemma 11.4.4. *If $t = n^{-1/\varepsilon'(H)}$, and p is such that $\gamma \to \infty$ as $n \to \infty$, then almost every $G \in \mathcal{G}(n,p)$ lies in $\mathcal{P}_H$.*

Proof. By our new choice of t and definition of γ we now have

$$p = \gamma n^{-1/\varepsilon'}, \tag{4}$$

where $\varepsilon' := \varepsilon'(H)$. ε'

Before we start on the main proof, let us note an inequality for later use. For all $n \geqslant k$,

$$\begin{aligned} \binom{n}{k} n^{-k} &= \frac{1}{k!}\left(\frac{n}{n}\cdots\frac{n-k+1}{n}\right) \\ &\geqslant \frac{1}{k!}\left(\frac{n-k+1}{n}\right)^{k} \\ &= \frac{1}{k!}\left(1-\frac{k-1}{k}\right)^{k}. \end{aligned} \tag{5}$$

When n gets large and k is bounded, as in our case, the upshot is that n^k exceeds $\binom{n}{k}$ by no more than a constant factor, one independent of n.

Our goal is to apply Lemma 11.4.2, and hence to bound $\sigma^2/\mu^2 = \big(\mathbb{E}(X^2) - \mu^2\big)/\mu^2$ from above. To help us estimate $\mathbb{E}(X^2)$, we begin by rewriting X^2 in a strange way, as

$$X^2(G) = \big|\{H \in \mathcal{H} : H \subseteq G\}\big|^2 = \big|\{(H', H'') \in \mathcal{H}^2 : H' \subseteq G \text{ and } H'' \subseteq G\}\big|.$$

We can now calculate $\mathbb{E}(X^2)$ by double counting, just as in (1):

$$\mathbb{E}(X^2) = \sum_{(H',H'')\in\mathcal{H}^2} \mathbb{P}\,[\,H' \cup H'' \subseteq G\,]\,. \tag{6}$$

Let us then calculate these probabilities $\mathbb{P}\,[\,H' \cup H'' \subseteq G\,]$. Given $H', H'' \in \mathcal{H}$, we have

$$\mathbb{P}\,[\,H' \cup H'' \subseteq G\,] \;=\; p^{2\ell - \|H'\cap H''\|}\,.$$

i As $\|H' \cap H''\| \leqslant i\varepsilon'$ for $i := |H' \cap H''|$ by definition of ε', this yields

$$\mathbb{P}\,[\,H' \cup H'' \subseteq G\,] \;\leqslant\; p^{2\ell - i\varepsilon'}\,. \tag{7}$$

We have now estimated the individual summands in (6); what does this imply for the sum as a whole? Since (7) depends on the parameter $i = |H' \cap H''|$, we partition the range $\mathcal{H}^2$ of the sum in (6) into the subsets

$\mathcal{H}^2_i$
$$\mathcal{H}^2_i := \{\, (H', H'') \in \mathcal{H}^2 : |H' \cap H''| = i \,\}\,, \qquad i = 0, \dots, k,$$

and calculate for each $\mathcal{H}^2_i$ the corresponding sum

A_i
$$A_i := \sum\nolimits_i \mathbb{P}\,[\,H' \cup H'' \subseteq G\,]$$

$\sum_i$ by itself. (Here, as below, we use $\sum_i$ to denote sums over all pairs $(H', H'') \in \mathcal{H}^2_i$.)

If $i = 0$ then H' and H'' are disjoint, so the events $H' \subseteq G$ and $H'' \subseteq G$ are independent. Hence,

$$\begin{aligned}
A_0 &= \sum\nolimits_0 \mathbb{P}\,[\,H' \cup H'' \subseteq G\,] \\
&= \sum\nolimits_0 \mathbb{P}\,[\,H' \subseteq G\,] \cdot \mathbb{P}\,[\,H'' \subseteq G\,] \\
&\leqslant \sum_{(H',H'')\in\mathcal{H}^2} \mathbb{P}\,[\,H' \subseteq G\,] \cdot \mathbb{P}\,[\,H'' \subseteq G\,] \\
&= \Big(\sum_{H'\in\mathcal{H}} \mathbb{P}\,[\,H' \subseteq G\,]\Big) \cdot \Big(\sum_{H''\in\mathcal{H}} \mathbb{P}\,[\,H'' \subseteq G\,]\Big) \\
&\underset{(1)}{=} \mu^2.
\end{aligned} \tag{8}$$

Let us now estimate A_i for $i \geqslant 1$. Note that $\sum_i$ can be written as $\sum_{H'\in\mathcal{H}} \sum_{H''\in\mathcal{H}:\,|H'\cap H''|=i}$. For fixed H', the second sum ranges over

$$\binom{k}{i}\binom{n-k}{k-i} h$$

summands: the number of graphs $H'' \in \mathcal{H}$ with $|H'' \cap H'| = i$. Hence, for all $i \geqslant 1$ and suitable constants c_1, c_2 independent of n,

$$\begin{aligned}
A_i \;&=\; \sum\nolimits_i \mathbb{P}\,[\,H' \cup H'' \subseteq G\,] \\
&\underset{(7)}{\leqslant} \sum_{H' \in \mathcal{H}} \binom{k}{i}\binom{n-k}{k-i} h\, p^{2\ell} p^{-i\varepsilon'} \\
&\underset{(4)}{=} |\mathcal{H}| \binom{k}{i}\binom{n-k}{k-i} h\, p^{2\ell} \left(\gamma\, n^{-1/\varepsilon'}\right)^{-i\varepsilon'} \\
&\leqslant |\mathcal{H}|\, p^{\ell} c_1\, n^{k-i}\, h\, p^{\ell} \gamma^{-i\varepsilon'}\, n^{i} \\
&\underset{(1,2)}{=} \mu\, c_1 n^k h\, p^{\ell} \gamma^{-i\varepsilon'} \\
&\underset{(5)}{\leqslant} \mu\, c_2 \binom{n}{k} h\, p^{\ell} \gamma^{-i\varepsilon'} \\
&\underset{(1,2,3)}{=} \mu^2 c_2 \gamma^{-i\varepsilon'} \\
&\leqslant \mu^2 c_2 \gamma^{-\varepsilon'}
\end{aligned}$$

if $\gamma \geqslant 1$. By definition of the A_i, this implies with $c_3 := k c_2$ that

$$\mathbb{E}(X^2)/\mu^2 \underset{(6)}{=} \Big(A_0/\mu^2 + \sum_{i=1}^{k} A_i/\mu^2\Big) \underset{(8)}{\leqslant} 1 + c_3 \gamma^{-\varepsilon'}$$

and hence

$$\frac{\sigma^2}{\mu^2} = \frac{\mathbb{E}(X^2) - \mu^2}{\mu^2} \leqslant c_3 \gamma^{-\varepsilon'} \underset{\gamma \to \infty}{\longrightarrow} 0\,,$$

since $\varepsilon' \geqslant \varepsilon > 0$ by our assumption that $\ell = \|H\| > 0$. By Lemma 11.4.2, therefore, $X > 0$ almost surely, i.e. almost every $G \in \mathcal{G}(n,p)$ has a subgraph isomorphic to H and hence lies in $\mathcal{P}_H$. □

Theorem 11.4.5. (Erdős & Rényi 1960; Bollobás 1981)
Let H be a graph with at least one edge. Then $t = n^{-1/\varepsilon'(H)}$ is a threshold function for $\mathcal{P}_H$.

Proof. We have to show that almost no $G \in \mathcal{G}(n,p)$ lies in $\mathcal{P}_H$ if $\gamma \to 0$ as $n \to \infty$, and that almost all $G \in \mathcal{G}(n,p)$ lie in $\mathcal{P}_H$ if $\gamma \to \infty$ as $n \to \infty$. This latter assertion was proved in Lemma 11.4.4.

To prove that almost no $G \in \mathcal{G}(n,p)$ lies in $\mathcal{P}_H$ if $\gamma \to 0$, we apply Lemma 11.4.3 to a subgraph $H' \subseteq H$ for which the maximum in the definition of $\varepsilon'(H)$ is attained, i.e. which is such that $\varepsilon(H') = \varepsilon'(H)$. The lemma implies that almost no $G \in \mathcal{G}(n,p)$ contains a copy of H'. Since any graph containing H also contains H', this implies that almost no $G \in \mathcal{G}(n,p)$ contains a copy of H, as desired. □

The bound in Theorem 11.4.5 is particularly easy to compute for *balanced* graphs H, those for which $\varepsilon'(H) = \varepsilon(H)$. Cycles and trees are examples of balanced graphs. For cycles, we have the threshold familiar from the proof of Erdős's theorem:

Corollary 11.4.6. *If $k \geqslant 3$, then $t(n) = n^{-1}$ is a threshold function for the property of containing a k-cycle.* □

Note that t does not depend on k. (See also Exercise 18.)

For trees, there is a similar phenomenon. Here, the threshold function does depend on the order of the tree, but not on its shape:

Corollary 11.4.7. *If T is a tree of order $k \geqslant 2$, then $t(n) = n^{-k/(k-1)}$ is a threshold function for the property of containing a copy of T.* □

The systematic study of threshold functions has led to an overall picture of how the typical properties of a graph $G \in \mathcal{G}(n,p)$ unfold as the growth rate of $p = p(n)$ increases. This picture, dubbed the *evolution of random graphs*, is quite fascinating: as in the evolution of species, changes happen 'in jumps', marked by the times the growth rate of p crosses that of a threshold function.

For a very rough sketch, let us begin with edge probabilities p whose order of magnitude lies below n^{-2}; for such p, a random graph $G \in \mathcal{G}(n,p)$ almost surely has no edges at all. But as p grows, it acquires more and more structure. As p approaches n^{-1}, its components become larger and larger trees (Corollary 11.4.7), until at $p = n^{-1}$ the first cycles are born (Exercise 18). Soon, some of these will have several crossing chords, making the graph non-planar. At the same time, one component outgrows the others, until it devours them around $p = (\log n)n^{-1}$, making the graph connected. Hardly later, at a mere $p = (1+\epsilon)(\log n)n^{-1}$, our random graph already almost surely has a Hamilton cycle...

Exercises

1.$^-$ What is the probability that a random graph in $\mathcal{G}(n,p)$ has exactly m edges, for $0 \leqslant m \leqslant \binom{n}{2}$ fixed?

2. What is the expected number of edges in $G \in \mathcal{G}(n,p)$?

3. What is the expected number of K^r-subgraphs in $G \in \mathcal{G}(n,p)$?

4. Characterize the graphs that occur as a subgraph in every graph of sufficiently large average degree.

5. In the usual terminology of measure spaces (and in particular, of probability spaces), the phrase 'almost all' is used to refer to a set of points whose complement has measure zero. Rather than considering a limit of probabilities in $\mathcal{G}(n,p)$ as $n \to \infty$, would it not be more natural to define a probability space on the set of *all* finite graphs (one copy of each) and to investigate properties of 'almost all' graphs in this space, in the sense above?

6. Show that if almost all $G \in \mathcal{G}(n,p)$ have a graph property $\mathcal{P}_1$ and almost all $G \in \mathcal{G}(n,p)$ have a graph property $\mathcal{P}_2$, then almost all $G \in \mathcal{G}(n,p)$ have both properties, i.e. have the property $\mathcal{P}_1 \cap \mathcal{P}_2$.

7. Show that, for constant $p \in (0,1)$, almost every graph in $\mathcal{G}(n,p)$ has diameter 2.

8. Show that, for constant $p \in (0,1)$, almost no graph in $\mathcal{G}(n,p)$ has a separating complete subgraph.

9. Derive Proposition 11.3.1 from Lemma 11.3.2.

10. Show that for every graph H there exists a function $p = p(n)$ such that $\lim_{n\to\infty} p(n) = 0$ but almost every $G \in \mathcal{G}(n,p)$ contains an induced copy of H.

11.$^+$ (i) Show that, for every $0 < \epsilon \leqslant 1$ and $p = (1-\epsilon)(\ln n)n^{-1}$, almost every $G \in \mathcal{G}(n,p)$ has an isolated vertex.

(ii) Find a probability $p = p(n)$ such that almost every $G \in \mathcal{G}(n,p)$ is disconnected but the expected number of spanning trees of G tends to infinity as $n \to \infty$.

(Hint for (ii): A theorem of Cayley states that K^n has exactly n^{n-2} spanning trees.)

12.$^+$ Given $r \in \mathbb{N}$, find a $c > 0$ such that, for $p = cn^{-1}$, almost every $G \in \mathcal{G}(n,p)$ has a K^r minor. Can c be chosen independently of r?

13. Find an increasing graph property without a threshold function, and a property that is not increasing but has a threshold function.

14.$^-$ Let H be a graph of order k, and let h denote the number of graphs isomorphic to H on some fixed set of k elements. Show that $h \leqslant k!$. For which graphs H does equality hold?

15.$^-$ For every $k \geqslant 1$, find a threshold function for $\{\, G \mid \Delta(G) \geqslant k \,\}$.

16.$^-$ For every $d \in \mathbb{N}$, determine the threshold function for the property of containing the d-dimensional cube (see Exercise 2, Chapter 1), and for the property of containing the complete graph K^d.

17. Does the property of containing *any* tree of order k (for $k \geqslant 2$ fixed) have a threshold function? If so, which? If not, why not?

18. Show that $t(n) = n^{-1}$ is also a threshold function for the property of containing any cycle.

19. Consider the terms A_0 and A_1 in the proof of Lemma 11.4.4, which are both functions of n. Recall that $\mathbb{P}[\,H' \cup H'' \subseteq G\,] = p^{2\ell}$ both for $H' \cap H'' = \emptyset$ and for $|H' \cap H''| = 1$.
    (i) Show that $A_0 \not\to 0$ (while $A_1 \to 0$) as $n \to \infty$.
    (ii) Explain the difference without doing any formal calculations.

20.$^+$ Given a graph H, let $\mathcal{P}$ be the property of containing an induced copy of H. Show that $\mathcal{P}$ has no threshold function unless H is complete.

Notes

There are a number of monographs and texts on the subject of random graphs. The first comprehensive monograph was B. Bollobás, *Random Graphs*, Academic Press 1985. Another advanced monograph is S. Janson, T. Łuczak & A. Ruciński, *Random Graphs*, Wiley 2000; this concentrates on areas developed since *Random Graphs* was published. E.M. Palmer, *Graphical Evolution*, Wiley 1985, covers material similar to parts of *Random Graphs* but is written in a more elementary way. Compact introductions going beyond what is covered in this chapter are given by B. Bollobás, *Modern Graph Theory*, Springer GTM 184, 1998, and by M. Karoński, *Handbook of Combinatorics* (R.L. Graham, M. Grötschel & L. Lovász, eds.), North-Holland 1995.

A stimulating advanced introduction to the use of random techniques in discrete mathematics more generally is given by N. Alon & J.H. Spencer, *The Probabilistic Method*, Wiley 1992. One of the attractions of this book lies in the way it shows probabilistic methods to be relevant in proofs of entirely deterministic theorems, where nobody would suspect it. Other examples for this phenomenon are Alon's proof of Theorem 5.4.1, or the proof of Theorem 1.3.4; see the notes for Chapters 5 and 1, respectively.

The probabilistic method had its first origins in the 1940s, one of its earliest results being Erdős's probabilistic lower bound for Ramsey numbers (Theorem 11.1.3). Lemma 11.3.2 about the properties $\mathcal{P}_{i,j}$ is taken from Bollobás's Springer text cited above. A very readable rendering of the proof that, for constant p, every first order sentence about graphs is either almost surely true or almost surely false, is given by P. Winkler, Random structures and zero-one laws, in (N.W. Sauer et al., eds.) *Finite and Infinite Combinatorics in Sets and Logic* (NATO ASI Series C **411**), Kluwer 1993.

Theorem 11.3.5 is due to P. Erdős and A. Rényi, Asymmetric graphs, *Acta Math. Acad. Sci. Hungar.* **14** (1963), 295–315. For further references about the infinite random graph R see the notes in Chapter 8.

The seminal paper on graph evolution is P. Erdős & A. Rényi, On the evolution of random graphs, *Publ. Math. Inst. Hungar. Acad. Sci.* **5** (1960), 17–61. This paper also includes Theorem 11.4.5 for balanced graphs. The generalization to unbalanced subgraphs was first proved by Bollobás in 1981; see Karoński's *Handbook* chapter. The fact that all 'non-trivial' increasing graph properties have a threshold function was proved by B. Bollobás and A.G. Thomason, Threshold functions, *Combinatorica* **7** (1987), 35–38.

There is another way of defining a random graph G, which is just as natural and common as the model we considered. Rather than choosing the edges of G independently, we choose the entire graph G uniformly at random from among all the graphs on $\{0, \ldots, n-1\}$ that have exactly $M = M(n)$ edges: then each of these graphs occurs with the same probability of $\binom{N}{M}^{-1}$, where $N := \binom{n}{2}$. Just as we studied the likely properties of the graphs in $\mathcal{G}(n, p)$ for different functions $p = p(n)$, we may investigate how the properties of G in the other model depend on the function $M(n)$. If M is close to pN, the expected number of edges of a graph in $\mathcal{G}(n, p)$, then the two models behave very similarly. It is then largely a matter of convenience which of them to consider; see Bollobás for details.

In order to study threshold phenomena in more detail, one often considers the following *random graph process*: starting with a $\overline{K^n}$ as stage zero, one chooses additional edges one by one (uniformly at random) until the graph is complete. This is a simple example of a Markov chain, whose Mth stage corresponds to the 'uniform' random graph model described above. A survey about threshold phenomena in this setting is given by T. Łuczak, The phase transition in a random graph, in (D. Miklós, V.T. Sós & T. Szőnyi, eds.) *Paul Erdős is 80*, Vol. 2, Proc. Colloq. Math. Soc. János Bolyai (1996).