Four precious jewels
1. The random graph

Peter J. Cameron
University of St Andrews

Hangzhou, August 2019

The lectures

In the course of these lectures, I will speak about four remarkable objects:

- the Erdős–Rényi random graph, aka the Rado graph;
- the rational numbers (as ordered set);
- the Urysohn metric space;
- the pseudo-arc.

Although these objects are four individuals, there are various general theories that connect them, and I will speak about some of these:

- countably categorical structures;
- homogeneous structures and Fraïssé's Theorem;
- Ramsey classes and extreme amenability;
- inverse limits and the dual of Fraïssé's Theorem.

Finite random graphs

I think that highly symmetric objects are the most interesting, like the Petersen graph:

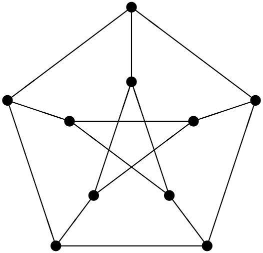

We choose a random graph on a given vertex set as follows: for each pair or vertices, decide independently whether to join them by an edge or not (for example, by the toss of a coin).

It turns out that the probability that the random graph on n vertices has any non-trivial automorphisms tends very rapidly to zero as $n \to \infty$.

So beautiful finite objects like the Petersen graph are rare …

The countable random graph

But a “phase shift” occurs when we go to the countably infinite.
In 1963, Erdős and Rényi showed that the random graph on a countable set of vertices has infinitely many automorphisms with probability 1.
The reason is even more extraordinary:

Theorem

There exists a graph R with the property that, if a countable random graph X is chosen at random, then with probability 1, X is isomorphic to R.

So there is only one countable random graph!
The graph R is the first of the four “precious jewels” of my title.
I will show you the proof.

Graphs and induced subgraphs

A graph consists of a set of vertices and a set of edges joining pairs of vertices; no loops, multiple edges, or directed edges are allowed.

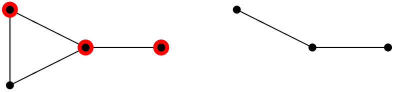

An induced subgraph of a graph consists of a subset of the vertex set together with all edges contained in the subset. In other words we are not allowed to delete edges within our chosen vertex set.

Alice's restaurant

The proof depends on the following property:

> Given two finite disjoint sets U and V of vertices, there is a vertex z which is joined to every vertex in U and to no vertex in V.

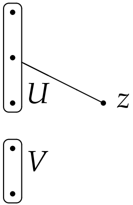

The point z is called a witness for the sets U and V. This is called the Alice's Restaurant property, or AR for short, after the song by Arlo Guthrie:

> *You can get anything you want*
> *At Alice's restaurant.*

The proof

I will show two things, which together give the proof.

Fact 1. With probability 1, a random countable graph has AR.

Fact 2. Any two countable graphs satisfying AR are isomorphic.

To prove Fact 1, we use from measure theory the fact that a countable union of null sets is null. Since there are only countably many choices for the (finite disjoint) sets U and V, it suffices to show that for a fixed choice of U and V the probability that no witness z exists is 0.

Suppose that $|U \cup V| = n$. Then the probability that a given vertex z is not the required witness is $1 - \frac{1}{2^n}$.

Since all edges are independent, the probability that none of $z_1, z_2, \ldots, z_N$ is the required witness is $\left(1 - \frac{1}{2^n}\right)^N$, which tends to 0 as $N \to \infty$.

So the event that no witness exists has probability 0, as required.

Proof of Fact 2

We use a method known to logicians as “back and forth”. Suppose that Γ_1 and Γ_2 are countable graphs satisfying AR: enumerate their vertex sets as $(a_0, a_1, \ldots)$ and $(b_0, b_1, \ldots)$. We build an isomorphism ϕ between them in stages.

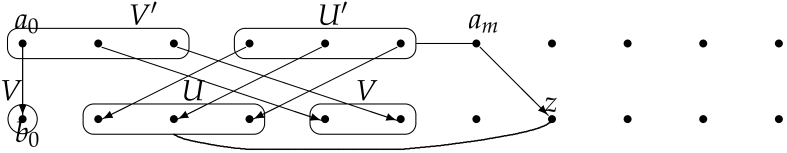

At stage 0, map a_0 to b_0.

At even-numbered stages, let a_m the first unmapped a_i. Let U' and V' be its neighbours and non-neighbours among the vertices alreay mapped, and let U and V be their images under ϕ. Use AR in graph Γ_2 to find a witness z for U and V. Then map a_m to z.

Fact 2, continued

At odd-numbered stages, go in the other direction, using AR in Γ_1 to choose a pre-image of the first unmapped vertex in Γ_2. This approach guarantees that every vertex of Γ_1 occurs in the domain, and every vertex of Γ_2 in the range, of ϕ; so we have constructed an isomorphism.

The proof is finished. This is a fine example of a non-constructive existence proof: if almost all graphs have the property, then certainly a graph with the property exists. Erdős and Rényi didn't bother with an explicit construction.

Had we only gone "forward", we would only use property AR in Γ_2, and we would have constructed an embedding, but could not guarantee that it is onto.

Properties of R

Recall that a countable graph Γ is universal if every finite or countable graph can be embedded into Γ as induced subgraph.

Fact 3. R is universal (for finite and countable graphs).

To see this, revisit the back-and-forth "machine" but use it only in the forward direction. As we saw, this only requires AR to hold in Γ_2, and delivers an embedding of Γ_1 in Γ_2.

A graph Γ is homogeneous if every isomorphism between finite induced subgraphs of Γ can be extended to an automorphism of Γ. (This is a very strong symmetry condition: it immediately implies that Γ is vertex-transitive, arc-transitive, etc.)

Fact 4. R is homogeneous.

To see this, take $\Gamma_1 = \Gamma_2 = R$, and start the back-and-forth machine from the given finite isomorphism ϕ_0; the result is an automorphism of R extending ϕ_0.

Rado's construction

In 1964, about the same time as Erdős and Rényi, Richard Rado gave the following construction. (I don't think that Rado knew what Erdős and Rényi were doing, or *vice versa*.)

The vertex set of Rado's graph R is the set $\mathbb{N}$ of natural numbers (including 0).

Given two vertices x and y, with $x < y$, we join x to y if, when y is written in base 2, its x-th digit is 1 – in other words, if we write y as a sum of distinct powers of 2, one of these is 2^x.

Don't forget that the graph is undirected! Thus

- 0 is joined to all odd numbers;
- 1 is joined to 0 and to all numbers congruent to 2 or 3 (mod 4).
- $\ldots$

Rado's graph is indeed an example of the random graph. To prove this, we have to verify AR. Given U and V, choose $m > \max\{U \cup V\}$ and put $z = \sum_{u \in U} 2^u + 2^m$.

A number-theoretic construction

Since the prime numbers are “random”, we should be able to use them to construct the random graph. Here’s how.

Recall that, if p is an odd prime not dividing a, then a is a quadratic residue (mod p) if the congruence $a \equiv x^2 \pmod{p}$ has a solution, and a quadratic non-residue otherwise. A special case of the law of quadratic reciprocity, due to Gauss, asserts that if the primes p and q are congruent to 1 (mod 4), then p is a quadratic residue (mod q) if and only if q is a quadratic residue (mod p).

So we can construct a graph whose vertices are all the prime numbers congruent to 1 (mod 4), with p and q joined if and only if p is a quadratic residue (mod q): the law of quadratic reciprocity guarantees that the edges are undirected.

This graph is isomorphic to the random graph!
To show this we have to verify AR. So let U and V be finite disjoint sets of primes congruent to 1 (mod 4). For each $u_i \in U$ let a_i be a fixed quadratic residue (mod u_i); for each $v_j \in V$, let b_j be a fixed quadratic non-residue mod v_j.
By the Chinese Remainder Theorem, the simultaneous congruences

- $z \equiv a_i \pmod{u_i}$ for all $u_i \in U$,
- $z \equiv b_j \pmod{v_j}$ for all $v_j \in V$,
- $z \equiv 1 \pmod{4}$,

have a solution modulo $4 \prod u_i \prod v_j$. By Dirichlet's Theorem, this congruence class contains a prime, which is the required witness.

The Skolem paradox

The downward Löwenheim–Skolem theorem of model theory says that a consistent theory in a countable first-order language has a countable model. (More about first-order logic coming up soon.)

The Skolem paradox is this: There is a theorem of set theory (for example, as axiomatised by the Zermelo–Fraenkel axioms) which asserts the existence of uncountable sets. Assuming that ZF is consistent (as we all believe!), how can this theory have a countable model?

My point here is not to resolve this paradox, but to use it constructively.

A set-theoretic construction

Let M be a countable model of the Zermelo–Fraenkel axioms for set theory. Then M consists of a collection of things called "sets", with a single binary relation $\in$, the "membership relation".

Form a graph on the set M by joining x and y if either $x \in y$ or $y \in x$.

This graph turns out to be the random graph!

Indeed, the precise form of the axioms is not so important. We need a few basic axioms (Empty Set, Pairing, Union) and, crucially, the Axiom of Foundation, and that is all. It does not matter, for example, whether or not the Axiom of Choice holds. (The Axiom of Foundation forbids infinite descending chains $\cdots x_2 \in x_1 \in x_0$ under the membership relation.)

Back to Rado's graph

In the set-theoretic construction, it doesn't matter whether the Axiom of Infinity holds or not.

There is a simple description of a model of set theory in which the negation of the axiom of infinity holds (called hereditarily finite set theory). In this theory, all sets are finite, all their subsets are finite, and so on.

We represent sets by natural numbers. We encode a finite set $\{a_1, \ldots, a_r\}$ of natural numbers by the natural number $2^{a_1} + \cdots + 2^{a_r}$. (So, for example, 0 encodes the empty set.)

When we apply the construction of "symmetrising the membership relation" to this model, we obtain exactly Rado's description of his graph!

Group-theoretic properties

Here are some properties of the graph R and its automorphism group.

- $\mathrm{Aut}(R)$ has cardinality $2^{\aleph_0}$.
- $\mathrm{Aut}(R)$ is simple.
- $\mathrm{Aut}(R)$ has the strong small index property: this means that any subgroup of this group with index strictly smaller than $2^{\aleph_0}$ lies between the pointwise and setwise stabilisers of a finite set.
- As a consequence, any graph Γ on fewer than $2^{\aleph_0}$ vertices satisfying $\mathrm{Aut}(\Gamma) \cong \mathrm{Aut}(R)$ is isomorphic to R.
- All cycle structures of automorphisms of R are known.
- R is a Cayley graph for a wide class of countable groups, including all countable abelian groups of infinite exponent. For these groups, a "random Cayley graph" is isomorphic to R with probability 1.

First-order logic

My treatment of first-order logic will be rather brief. It describes structures, one of which is a set (called the domain) equipped with functions or operations, relations, and constants; typical examples are groups, graphs, ordered fields, ...

Each function, relation or constant is represented by a symbol in the language of the logic, which also has variables, connectives, and quantifiers. We are only allowed finite conjunctions, disjunctions and quantifications, and we are only allowed to quantify over the domain, not over subsets or functions. (Higher-order logics remove some of these restrictions.)

A variable is free if it is not quantified; a sentence is a formula with no free variables. Examples include

- $(\forall x)(\forall y)(\forall z)(x \circ (y \circ z)) = ((x \circ y) \circ z)$ (the associative law);
- $(\forall x)(\forall y)((x \sim y) \vee (\exists z)((x \sim z) \wedge (y \sim z)))$ (stating that a graph has diameter 2).

Categoricity

Imagine that we are trying to write sentences which are axioms for some theory. Can we specify a structure completely by first-order sentences?

It follows from the Löwenheim–Skolem theorems that, if an infinite structure satisfies a collection of sentences, then there are arbitrarily large structures satisfying them. So we cannot specify cardinality!

So we compromise as follows. If α is an infinite cardinal number, we say that a set of sentences is α-categorical if any two structures of cardinality α satisfying them are isomorphic.

A theorem of Morley asserts that, for countable first-order languages, there are only two kinds of cardinality, countable and uncountable: if α and β are uncountable cardinals, then α-categorical is equivalent to β-categorical.

If a set of sentences is α-categorical, we also describe the unique countable structure satisfying it as being α-categorical.

The Engeler–Ryll-Nardzewski–Svenonius Theorem

A permutation group G on an infinite set Ω is oligomorphic if it has only finitely many orbits on Ω^n for all n. (This is a very strong symmetry condition.)

The theorem of Engeler, Ryll-Nardzewski and Svenonius (proved independently by them in 1959) asserts a remarkable equivalence between categoricity and symmetry:

Theorem

A countable first-order structure is countably categorical if and only if its automorphism group is oligomorphic.

The random graph

The random graph is countably categorical. The Alice's Restaurant property which characterises it can be expressed as a set of first-order sentences. For example, the case $|U| = 2$, $|V| = 1$ can be written (slightly abbreviated) as

$$\begin{aligned}(\forall u_1)(\forall u_2)(\forall v) \quad & ((u_1 \neq v \neq u_2 \neq u_1) \Rightarrow \\ & (\exists z)(z \sim u_1) \wedge (z \sim u_2) \wedge (z \not\sim v)).\end{aligned}$$

The automorphism group of the random graph is oligomorphic. For two n-tuples of distinct points lie in the same orbit if and only if the (labelled) subgraphs they induce are isomorphic. So the number of orbits on such tuples is equal to $2^{n(n-1)/2}$, the number of labelled graphs on n vertices.
So the random graph illustrates the theorem.