Lecture 3

What did we do so far?

- We defined **density function** of graph properties.
- We defined **asymptotic probabilities** of graph properties.
- We introduced **Second Order Logic** and some of its fragments.
- We studied **many examples**.
- We stated **two theorems**:

  A: the 0-1 law (GKLT 1969)
  for FOL-definable graph properties and

  B: the modular periodicity theorem (Specker-Blatter 1981)
  for $MSOL$-definable graph properties.

In lecture 2 we gave the background in logic.

0-1 Laws

We want to **prove** in this lecture

Theorem 1
Let $\mathcal{P}$ be a FOL-definable graph property and $p = \frac{1}{2}$ be constant. Then the asymptotic probability $\mu_{\mathcal{P}} = \lim_{n \to \infty} \mu_{\mathcal{P}}(n)$ always exists and $\mu_{\mathcal{P}} = 0$ or $\mu_{\mathcal{P}} = 1$.

Remarks on Theorem 1

More generally, let $\mathcal{K}$ be a class of τ-structures closed under isomorphisms. We say that $\mathcal{K}$ has a 0-1-law for **FOL**(τ) if for every $\phi \in$ **FOL**(τ) the fraction of τ-structures in $\mathcal{K}$ which satisfy ϕ tends either to 0 or to 1.

The fraction and probabilities here are computed by dividing by $f_{\mathcal{K}(n)}$ rather than by $gr(n)$.

Note: In Theorem 1 we consider only **graph properties**,
i.e. $\mathcal{K} = Graphs$ consists of all finite structures with
one symmetric, irreflexive binary relation.

Exercise: Formulate Theorem 1 for **any** τ
and the $\mathcal{K}$ class of **all finite** τ**-structures**.

Exercise: Show that with respect to all structures with one binary relation $\mu(Graphs) = 0$.

Exercise: Modify the proof of Theorem 1
such that it works for $\mathcal{K}$ consisting of all finite τ-structures, including, say, **directed graphs**.

Exercise: Find a **FOL**(τ)-definable $\mathcal{K}$ which has no 0-1-law for **FOL**(τ).

The theory $T^{\tau}_{p(n)}$

for $\mathcal{K}$ consisting of all simple graphs

Definition 2

We de note by $T^{\tau}_{p(n)}$ the set $T^{\tau}_{p(n)} = \{\phi \in FOL(\tau) : \mu_{\phi} = 1\}$.

Proposition 3

Let $p(n)$ be the probability the edge probability between two among n vertices.

(i) $T^{\tau}_{p(n)}$ *is closed under logical consequence, i.e. if* $T^{\tau}_{p(n)} \models \phi$ *then* $\phi \in T^{\tau}_{p(n)}$.

(ii) $T^{\tau}_{p(n)}$ *is deductively closed, i.e. if* $T^{\tau}_{p(n)} \vdash \phi$ *then* $\phi \in T^{\tau}_{p(n)}$.

(iii) $T^{\tau}_{p(n)}$ *is consistent, i.e.* $\mathbf{FALSE} \notin T^{\tau}_{p(n)}$, *hence satisfiable.*

(iv) $T^{\tau}_{p(n)}$ *has no finite models.*

(v) $FOL(\tau)$ *with probability* $p(n)$ *satisfies the 0-1 Law iff* $T^{\tau}_{p(n)}$ *is complete, i.e. for all* $\phi \in FOL(\tau)$ *either* $\phi \in T^{\tau}_{p(n)}$ *or* $\neg\phi \in T^{\tau}_{p(n)}$

Useful facts about $\mu_\phi(n)$.

Proposition 4

Let ϕ and ψ be $FOL(\tau)$-sentences.

(i) $\mu_{\neg\phi}(n) = 1 - \mu_\phi(n)$.

(ii) If $\phi \models \psi$ *then* $\mu_\phi(n) \leq \mu_\psi(n)$.

(iii) $\mu_{\phi\vee\psi}(n) \leq \mu_\phi(n) + \mu_\psi(n)$.

(iv) More generally, if $\Phi = \phi_1 \vee \ldots \vee \phi_n$ *then* $\mu_\Phi(n) \leq \sum_i^n \mu_{\phi_i}(n)$.

Proof of Proposition 3

More details on the blackboard (and **homework**).

(i) We use the compactness theorem for FOL and proposition 4.

Let $\phi \in \mathbf{FOL}(\tau)$ with $T^{\tau}_{p(n)} \models \phi$. We have to show $\phi \in T^{\tau}_{p(n)}$.

By compactness there are $\phi_1, \ldots, \phi_k \in T^{\tau}_{p(n)}$ such that $\{\phi_1, \ldots, \phi_k\} \models \phi$.

By Proposition 4(ii) $\mu_{\phi}(n) \geq \mu_{\phi_1 \wedge \ldots \wedge \phi_k}(n)$.

By Proposition 4(i) $\mu_{\phi_1 \wedge \ldots \wedge \phi_k}(n) = 1 - \mu_{\neg\phi_1 \vee \ldots \vee \neg\phi_k}(n)$

By Proposition 4(iv) $\mu_{\phi}(n) \geq 1 - \sum_{i=1}^{k} \mu_a \neg\phi_i(n)$.

Because $\phi_1, \ldots, \phi_k \in T^{\tau}_{p(n)}$ the limit $\lim_{n\to\infty} \sum_{i=1}^{k} \mu_a \neg\phi_i(n) = 0$.

Thus $\mu_{\phi} = 1$, and $\phi \in T^{\tau}_{p(n)}$.

(ii) To show that $T^{\tau}_{p(n)}$ is deductively closed
we use the completenss theorem for FOL.

Proof of Proposition 3 (continued)

(iii) We have to show that $T^{\tau}_{p(n)}$ is consistent.

$\mu_{\mathbf{FALSE}}(n) = 0$. Hence $\mu_{\mathbf{FALSE}} = 0$ and $\mathbf{FALSE} \notin T^{\tau}_{p(n)}$.

(iv) We have to show that $T^{\tau}_{p(n)}$ has no finite models.

Let χ_k be the sentence "*there are at least* k *distinct elements*".
$\mu_{\chi_k}(n) = 1$ for $n \geq k$.

So for all $k \in \mathbb{N}$ the formula $\chi_k \in T^{\tau}_{p(n)}$ and $T^{\tau}_{p(n)}$ has no finite models.

(v) Follows from the definition of $T^{\tau}_{p(n)}$, Proposition 4
and the definition of completeness.

This completes the proof. Q.E.D.

The extension axioms $EA_{n,m}$

Let S be a finite set of cardinality n and $T \subseteq S$ of cardinality m.
The extension axiom $EA_{n,m}$ says that there exists a vertex $z \notin S$ such that for every $x \in T$ there is an edge between z and x and that for every $x \in S - T$ there is no edge between z and x.

Formally

$$\forall x_1, \ldots, x_n \left(\left(\bigwedge_{i \neq j} x_i \neq x_j \right) \to \exists z \left(\bigwedge_{i}^{n} z \neq x_i \wedge \bigwedge_{i \leq m} E(z, x_i) \wedge \bigwedge_{i > m}^{n} \neg E(z, x_i) \right) \right)$$

We denote by EA_k the formula $E_{2k,k}$.

Useful facts about the extension axioms

Proposition 5

(i) For $m' \geq m, n' - m' \geq n - m$ we have $EA_{n',m'} \models EA_{n,m}$.

(ii) For every k $\mu_{EA_k} = 1$, and hence $\mu_{EA_{n,m}} = 1$, for every n and $m \leq n$.

(iii) Each EA_k has arbitrarily large finite models.

(iv) For each $n, m \leq n$ we have $EA_{n,m} \in T^{\tau}_{\frac{1}{2}}$.

Proof of Proposition 5.

Details on the blackboard.

(i) Trivial from the definition of $EA_{n,m}$.

(ii) Follows from the stronger Proposition 6 below.

(iii) From (ii).

(iv) Also from (ii).

For which $p(n)$ is $\mu_{EA_k}^{p(n)} = 1$?

Proposition 6 (Bol85, page 40ff.)

Assume $p(n)$ is such that every $\epsilon > 0$ we have

$$p(n) \cdot n^{\epsilon} \to \infty$$

and

$$(1 - p(n)) \cdot n^{\epsilon} \to \infty.$$

Then for every fixed k we have $\mu_{EA_k}^{p(n)} = 1$.

The condition is verified for

- $p(n)$ *constant.*
- $p(n) = \frac{1}{\log n}$

but not for $p(n) = \frac{1}{n}$

Proof of Proposition 6

For $T \subseteq S$ we say that

a vertex $z \notin S$ fits S and T

if z has edges to all vertices in T and none to $S - T$. Let T and $S - T$ be both of size k, and let $z \notin S$.

- The probability that z fits S and T is $p^k(1-p)^k$.
- The probability, that z does not fit S and T is $1 - p^k(1-p)^k$.
- There are $n - 2k$ many choices for z, hence the probability that no z fits S and T is
$$(1 - p^k(1-p)^k)^{n-2k}.$$
- We estimate this using $p(n) \cdot n^{\epsilon} \to \infty$ and $(1 - p(n)) \cdot n^{\epsilon} \to \infty$, and get (complete it!)
$$(1 - p^k(1-p)^k)^{n-2k} \leq exp\{-(n-2k)p^k(1-p)^k\} \leq exp\{-n^{\frac{1}{2}}\}.$$
- There are $\binom{n}{k}\binom{n-k}{k} \leq n^{2k}$ many choices for S and T, hence
$$\mu^{p(n)}_{\neg EA_k}(n) \leq n^{2k} exp\{-n^{\frac{1}{2}}\} = o(n^{-1}).$$

  Hence $\mu^{p(n)}_{\neg EA_k} = 0$ and $\mu^{p(n)}_{EA_k} = 1$.

Homework: Details for Bollobas (Proposition 6)

Completed by Eyal Rozenberg

We want to show that

$$(1 - p^k \cdot (1-p)^k)^{n-2k} \longrightarrow 0 \text{ for } n \to \infty$$

We use that $p^k(1-p)^k < 1$.

We put $f = \left(p^k(1-p)^k\right)^{-1}$ and $g = n - 2k$.

Now

$$(1 - p^k \cdot (1-p)^k)^{n-2k} = (1 - \frac{1}{f})^g = \left(1 - (\frac{1}{f})^f\right)^{\frac{g}{f}}$$

The theory EA

We define EA to be the set $\{EA_k : k \in \mathbb{N}\}$.

Here are some facts about EA, which all follow from Propositions 3,4, 5 and 6.

Proposition 7

(i) $EA \subseteq T^{\tau}_{p(n)}$ *for every* $p(n)$ *with* $p(n) \cdot n^{\epsilon} \to \infty$ *and* $(1 - p(n)) \cdot n^{\epsilon} \to \infty$.

(ii) EA *is satisfiable.*

(iii) EA *has no finite models.*

EA is $\aleph_0$-categorical

We say that a set $T \subseteq FOL(\tau)$, a theory, is κ-**categorical** if T is satisfiable and all its models of cardinality κ are τ-isomorphic.

Proposition 8

EA is $\aleph_0$-categorical.

Proof: Details on the blackboard.

Use Cantor's method.

Vaught's Test

R. Vaught noticed the following useful connection between categoricity, completeness and decidability.

Proposition 9 (R. Vaught, 1954)

Let τ be countable and $T \subseteq FOL(\tau)$ κ-categorical for some infinite κ, and which has not finite models. Then

(i) T is complete.

(ii) If T is semi-computable, then the relation $T \models \phi$ is computable.
In other words, T is **decidable**.

Proof of Vaught's Test

(i) Assume neither $T \models \phi$ nor $T \models \neg\phi$.
So both $T \cup \{\phi\}$ and $T \cup \{\neg\phi\}$ are satisfiable, and have infinite models, and using compactness and the Löwenheim-Skolem Theorem, have models of cardinality κ.

Let $\mathfrak{A} \models T \cup \{\phi\}$ and $\mathfrak{B} \models T \cup \{\neg\phi\}$ of size κ.
As T is κ-categorical, $\mathfrak{A}$ and $\mathfrak{B}$ are isomorphic.
This contradicts $\mathfrak{A} \models T \cup \{\phi\}$ and $\mathfrak{B} \models T \cup \{\neg\phi\}$.

(ii) The consequences of a semi-computable theory are semi-computable, by Gödel's Completeness Theorem.

So both $\{\phi : T \models \phi\}$ and $\{\phi : T \models \neg\phi\}$ are semi-computable.
But $\{\phi : T \models \phi\} = FOL - \{\phi : T \models \neg\phi\}$, hence they are computable.

0-1 Law for FOL

Theorem 10

Assume $p(n) \cdot n^{\epsilon} \to \infty$ *and* $(1 - p(n)) \cdot n^{\epsilon} \to \infty$*. Then we have*

(i) EA *is complete and* $EA = T^{\tau}_{p(n)}$*.*

(ii) FOL *satisfies the 0-1 Law.*

(iii) It is decidable whether $\mu_{\phi} = 0$ *or* $\mu_{\phi} = 1$*.*
(No statement on complexity)

E. Grandjean showed that computing μ_{ϕ} is **PSpace**-complete.

PROJECT: E. Grandjean, Complexity of the first–order theory of almost all structures, Information and Control, vol. 52 (1983), 180-204.

Homework

- Let $\mathcal{K}$ be with $\mu(\mathcal{K}) = 1$. Show that $\mathcal{K}$ has a $0 - 1$-law for **FOL**(τ).

- Show that the graph property of being connected $CONN$ has $\mu(CONN) = 1$ on graphs.

- Show that the 0-1-law for graph properties restricted to connected graphs holds.

How to prove Theorem A?

- We have to define the logic $\mathcal{L}^{\omega}_{\infty,\omega}$, which has finitely many variables but infinite conjunctions and disjunctions.

- We have to define Ehrenfeucht-Fraïssé Games with pebbles.

- Then we show that

  For finite graphs G_1 and G_2 for which $EA_{n,m}$ holds for all $m \leq n \leq k$ we have that player **II** has a k-pebble winning strategy for the infinite game on G_1 and G_2.

  This replaces completeness of EA.