The convergence of three notions of limit for finite structures

Alex Kruckman

Indiana University, Bloomington

Workshop on model theory of finite
and pseudofinite structures
& Logic seminar

University of Leeds
11 April, 2018

Outline

1. Three perspectives on the (Rado) random graph:
   - Random construction
   - Fraïssé limit
   - Zero-one law
2. A definition: Strongly pseudofinite theories
3. A sufficient condition: Total amalgamation classes
4. A tool: The Aldous–Hoover–Kallenberg representation
5. Consequences and questions (work in progress, joint with C. Hill)

The Erdős–Renyi construction

For each $n \in \omega$, build a graph with domain $[n] = \{0, \ldots, n-1\}$:

- For each pair $i < j$, flip a fair coin.
- Set iEj iff the coin comes up heads.

This is the Erdős–Renyi process $G(n, 1/2)$.

The Erdős–Renyi construction

For each $n \in \omega$, build a graph with domain $[n] = \{0, \ldots, n-1\}$:

- For each pair $i < j$, flip a fair coin.
- Set iEj iff the coin comes up heads.

This is the Erdős–Renyi process $G(n, 1/2)$.

Let $\mathcal{G}(n)$ be the set of all graphs with domain $[n]$.

We obtain each graph with probability $2^{-\binom{n}{2}}$.
So $G(n, 1/2)$ corresponds to the uniform measure on $\mathcal{G}(n)$.

The random graph

There is also an infinite Erdős–Renyi process $G(\omega, 1/2)$: Flip countably many coins, one for each pair $i < j < \omega$.

$G(\omega, 1/2)$ builds a single graph up to isomorphism with probability 1: The (Rado) random graph $\mathcal{R}$.

The random graph

There is also an infinite Erdős–Renyi process $G(\omega, 1/2)$: Flip countably many coins, one for each pair $i < j < \omega$.

$G(\omega, 1/2)$ builds a single graph up to isomorphism with probability 1: The (Rado) random graph $\mathcal{R}$.

Extension property $E(A, B)$

For any two disjoint finite sets $A, B \subseteq \omega$, there is a vertex $c \in \omega$ such that cEa for all $a \in A$ and $\neg cEb$ for all $b \in B$.

Each instance $E(A, B)$ of the extension property is satisfied with probability 1 in $G(\omega, 1/2)$.

By a back-and-forth argument, $\mathcal{R}$ is the unique countable graph satisfying all the extension properties up to isomorphism.

The random graph

$\mathcal{R}$ also arises naturally in (at least) two other ways:

- $\mathcal{R}$ is the Fraïssé limit of the class of finite graphs.
- The class of finite graphs has a logical zero-one law (for the uniform measures), and $\mathcal{R}$ is the unique countable model for the limit theory.

Fraïssé classes

Conventions: L is always a finite relational language. I allow empty structures.

A *Fraïssé class* is a class K of finite L-structures, such that

1. K is closed under isomorphism.
2. K is closed under substructure (hereditary property).
3. K has the amalgamation property (2-amalgamation):

$$\begin{tikzcd} & D & \\ A \arrow[ur, dashed] & & B \arrow[ul, dashed] \\ & C \arrow[ul] \arrow[ur] & \end{tikzcd}$$

Fraïssé limits

Let K be a Fraïssé class.

There is a countable structure M_K, the *Fraïssé limit* of K, satisfying:

1. Universality: K is the class of finite substructures of M_K.
2. Homogeneity: Any isomorphism between finite substructures of M_K extends to an automorphism of M_K.

Moreover, M_K is unique up to isomorphism.

Fraïssé limits

Let K be a Fraïssé class.

There is a countable structure M_K, the *Fraïssé limit* of K, satisfying:

1. Universality: K is the class of finite substructures of M_K.
2. Homogeneity: Any isomorphism between finite substructures of M_K extends to an automorphism of M_K.

Moreover, M_K is unique up to isomorphism.

Let $T_K = \mathrm{Th}(M_K)$, the *generic theory* of K.

T_K is $\aleph_0$-categorical and has quantifier elimination.

Here is an axiomatization:

1. *Universal axioms.* For every finite structure $A \notin K$,
$$\forall \overline{x}\, \neg\theta_A(\overline{x}).$$
2. *Extension axioms.* For all $A \subseteq B$ in K with $|B| = |A| + 1$,
$$\forall \overline{x}\, \exists y\, (\theta_A(\overline{x}) \to \theta_B(\overline{x}, y)).$$

Here θ_C is the conjunction of the atomic diagram of the structure C.

The zero-one law for finite graphs

Let $\mu_n (= G(n, 1/2))$ be the uniform measure on $\mathcal{G}(n)$.

For any sentence φ, and any n, let $[\varphi]_{\mathcal{G}(n)} = \{G \in \mathcal{G}(n) \mid G \models \varphi\}$.

The zero-one law for finite graphs

Let $\mu_n (= G(n, 1/2))$ be the uniform measure on $\mathcal{G}(n)$.
For any sentence φ, and any n, let $[\varphi]_{\mathcal{G}(n)} = \{G \in \mathcal{G}(n) \mid G \models \varphi\}$.

Then for any $\varphi \in \mathrm{Th}(\mathcal{R}) = T_{\mathcal{G}}$,

$$\lim_{n\to\infty} \mu_n([\varphi]_{\mathcal{G}(n)}) = 1.$$

We say that $\mathrm{Th}(\mathcal{R})$ is the *almost-sure theory* of $(\mathcal{G}(n), \mu_n)_{n\in\omega}$.

The zero-one law for finite graphs

Let $\mu_n (= G(n, 1/2))$ be the uniform measure on $\mathcal{G}(n)$.
For any sentence φ, and any n, let $[\varphi]_{\mathcal{G}(n)} = \{G \in \mathcal{G}(n) \mid G \models \varphi\}$.

Then for any $\varphi \in \mathrm{Th}(\mathcal{R}) = T_{\mathcal{G}}$,

$$\lim_{n\to\infty} \mu_n([\varphi]_{\mathcal{G}(n)}) = 1.$$

We say that $\mathrm{Th}(\mathcal{R})$ is the *almost-sure theory* of $(\mathcal{G}(n), \mu_n)_{n\in\omega}$.

More generally, if $(X_n, \mu_n)_{n\in\omega}$ is any sequence such that μ_n is a probability measure on a space X_n of finite L-structures, we say that:

- $(\mu_n)_{n\in\omega}$ has a *zero-one law* if for every sentence φ,

  $$\lim_{n\to\infty} \mu_n([\varphi]_{X_n}) = 0 \text{ or } 1.$$

- If $(\mu_n)_{n\in\omega}$ has a zero-one law,

  $$T^{\text{a.s.}} = \{\varphi \mid \lim_{n\to\infty} \mu_n([\varphi]_{X_n}) = 1\}$$

  is the *almost-sure theory* of $(\mu_n)_{n\in\omega}$.

The case of linear orders

Generic theories and almost-sure theories do not agree in general.

The class $\mathcal{L}$ of finite linear orders is a Fraïssé class.
Fraïssé limit: $M_{\mathcal{L}} = (\mathbb{Q}, \leq)$.
Generic theory: $T_{\mathcal{L}} = \mathrm{DLO}$ (dense linear orders without endpoints).
Almost-sure theory: Infinite discrete linear orders with endpoints.

The case of linear orders

Generic theories and almost-sure theories do not agree in general.

The class $\mathcal{L}$ of finite linear orders is a Fraïssé class.
Fraïssé limit: $M_{\mathcal{L}} = (\mathbb{Q}, \leq)$.
Generic theory: $T_{\mathcal{L}} = \mathrm{DLO}$ (dense linear orders without endpoints).
Almost-sure theory: Infinite discrete linear orders with endpoints.

Definition

A theory T is *pseudofinite* if every sentence $\varphi \in T$ has a finite model.

- DLO is not pseudofinite: Consider $(\exists x \top) \wedge (\forall y \, \exists z \, (y < z))$.
- But any almost-sure theory is pseudofinite: every sentence has many finite models (in a sense measured by the μ_n).

The case of triangle-free graphs

The class $\mathcal{G}_\triangle$ of finite triangle-free graphs is a Fraïssé class.
Fraïssé limit: $M_{\mathcal{G}_\triangle} = \mathcal{H}$, the Henson graph.
Generic theory: $T_{\mathcal{G}_\triangle} = \mathrm{Th}(\mathcal{H})$.

The case of triangle-free graphs

The class $\mathcal{G}_\triangle$ of finite triangle-free graphs is a Fraïssé class.
Fraïssé limit: $M_{\mathcal{G}_\triangle} = \mathcal{H}$, the Henson graph.
Generic theory: $T_{\mathcal{G}_\triangle} = \mathrm{Th}(\mathcal{H})$.

Theorem (Kolaitis–Prömel–Rothschild)

The sequence $(\mu_n)_{n\in\omega}$ of uniform measures on $\mathcal{G}_\triangle(n)$ has a zero-one law. $T^{\mathrm{a.s.}}$ is the generic theory of bipartite graphs.

Hence $T^{\mathrm{a.s.}} \neq T_{\mathcal{G}_\triangle}$, e.g. since $\mathcal{H}$ contains cycles of length 5.

The case of triangle-free graphs

The class $\mathcal{G}_\triangle$ of finite triangle-free graphs is a Fraïssé class.
Fraïssé limit: $M_{\mathcal{G}_\triangle} = \mathcal{H}$, the Henson graph.
Generic theory: $T_{\mathcal{G}_\triangle} = \mathrm{Th}(\mathcal{H})$.

Theorem (Kolaitis–Prömel–Rothschild)

The sequence $(\mu_n)_{n\in\omega}$ of uniform measures on $\mathcal{G}_\triangle(n)$ has a zero-one law. $T^{\text{a.s.}}$ is the generic theory of bipartite graphs.

Hence $T^{\text{a.s.}} \neq T_{\mathcal{G}_\triangle}$, e.g. since $\mathcal{H}$ contains cycles of length 5.

So the uniform measures give the wrong answer. What about other sequences of measures?

Cherlin's question

Question (Cherlin)

Is the generic theory $T_{\mathcal{G}_\triangle}$ of triangle-free graphs pseudofinite?

Cherlin's question

Question (Cherlin)

Is the generic theory $T_{\mathcal{G}_\triangle}$ of triangle-free graphs pseudofinite?

This question appears to be very difficult!

For example, it is open whether there are finite triangle-free graphs satisfying the extension axioms over all base graphs of size 4.

Cherlin's question

Question (Cherlin)

Is the generic theory $T_{\mathcal{G}_\triangle}$ of triangle-free graphs pseudofinite?

This question appears to be very difficult!

For example, it is open whether there are finite triangle-free graphs satisfying the extension axioms over all base graphs of size 4.

It seems likely that for some $\varphi \in T_{\mathcal{G}_\triangle}$, the finite models of φ are sporadic:

- Only occur in certain sizes,
- Or must have a very regular structure,
- Or no finite models at all!

In contrast, for all $\varphi \in T_{\mathcal{G}}$, the finite models of φ are extremely common.

Making Cherlin's question easier

Question

Does $T_{\mathcal{G}_\triangle}$ arise as the almost-sure theory for some *reasonable* sequence of measures $(\mu_n)_{n\in\omega}$?

Making Cherlin's question easier

Question

Does $T_{\mathcal{G}_\triangle}$ arise as the almost-sure theory for some *reasonable* sequence of measures $(\mu_n)_{n\in\omega}$?

What does *reasonable* mean?

- Requiring the μ_n to be uniform measures on $\mathcal{G}_\triangle(n)$ is too strong.
- But we need some assumptions: We don't want to allow each μ_n to give measure 1 to a single graph G_n in some sporadic family.

In this talk, I will focus on one possible meaning of *reasonable*.

Coherent measures

$\mathrm{Str}_L(n)$ is the set of all L-structures with domain $[n]$.
For any formula $\varphi(\overline{x})$ and any tuple $\overline{a}$ from $[n]$, define

$$[\varphi(\overline{a})]_n = \{M \in \mathrm{Str}_L(n) \mid M \models \varphi(\overline{a})\}.$$

Definition

$(\mu_n)_{n\in\omega}$ is *coherent* if each μ_n is a probability measure on $\mathrm{Str}_L(n)$, and:

1. For all $\varphi(\overline{x})$ quantifier-free, all $\overline{a}$ from $[n]$, and all $n \leq m$, $\mu_n([\varphi(\overline{a})]_n) = \mu_m([\varphi(\overline{a})]_m)$.
2. For all $\varphi(\overline{x})$ quantifier-free, all $\overline{a}$ from $[n]$, and all $\sigma \in S_n$, $\mu_n([\varphi(\overline{a})]_n) = \mu_n([\varphi(\sigma(\overline{a}))]_n)$.
3. For all $\varphi(\overline{x})$ and $\psi(\overline{y})$ quantifier-free and $\overline{a}$ and $\overline{b}$ *disjoint* from $[n]$, $\mu_n([\varphi(\overline{a}) \wedge \psi(\overline{b})]_n) = \mu_n([\varphi(\overline{a})]_n)\mu_n([\psi(\overline{b})]_n)$.

Motivation: The Erdős–Renyi constructions $G(n, 1/2)$, which *cohere* to a random construction $G(\omega, 1/2)$ of countably infinite graphs.

Measures on $\mathrm{Str}_L(\omega)$

$\mathrm{Str}_L(\omega)$ is the space of all L-structures with domain ω.

The topology on $\mathrm{Str}_L(\omega)$ is generated by basic clopen sets of the form

$$[\varphi(\overline{a})] = \{M \in \mathrm{Str}_L(\omega) \mid M \models \varphi(\overline{a})\}$$

where $\varphi(\overline{x})$ is a quantifier-free formula and $\overline{a}$ is a tuple from ω.

Measures on $\mathrm{Str}_L(\omega)$

$\mathrm{Str}_L(\omega)$ is the space of all L-structures with domain ω.

The topology on $\mathrm{Str}_L(\omega)$ is generated by basic clopen sets of the form

$$[\varphi(\overline{a})] = \{M \in \mathrm{Str}_L(\omega) \mid M \models \varphi(\overline{a})\}$$

where $\varphi(\overline{x})$ is a quantifier-free formula and $\overline{a}$ is a tuple from ω.

A Borel probability measure on $\mathrm{Str}_L(\omega)$ is determined by the measure of each finite conjunction of atomic and negated atomic formulas (and we always have $\mu([a = b]) = 0$ when $a \neq b$).

In the case of the Erdős-Renyi construction, for example,

$$\mu\left(\left[\left(\bigwedge_{i=1}^{m} a_i E b_i\right) \wedge \left(\bigwedge_{j=1}^{n} \neg a_j E b_j\right)\right]\right) = \left(\frac{1}{2}\right)^{m+n}.$$

Measures on $\mathrm{Str}_L(\omega)$

Definition

$(\mu_n)_{n\in\omega}$ is *coherent* if each μ_n is a probability measure on $\mathrm{Str}_L(n)$, and:

1. For all $\varphi(\overline{x})$ quantifier-free, all $\overline{a}$ from $[n]$, and all $n \leq m$, $\mu_n([\varphi(\overline{a})]_n) = \mu_m([\varphi(\overline{a})]_m)$.
2. For all $\varphi(\overline{x})$ quantifier-free, all $\overline{a}$ from $[n]$, and all $\sigma \in S_n$, $\mu_n([\varphi(\overline{a})]_n) = \mu_n([\varphi(\sigma(\overline{a}))]_n)$.
3. For all $\varphi(\overline{x})$ and $\psi(\overline{y})$ quantifier-free and $\overline{a}$ and $\overline{b}$ *disjoint* from $[n]$, $\mu_n([\varphi(\overline{a}) \wedge \psi(\overline{b})]_n) = \mu_n([\varphi(\overline{a})]_n)\mu_n([\psi(\overline{b})]_n)$.

By condition (1), a coherent sequence induces a Borel probability measure on $\mathrm{Str}_L(\omega)$.

Example: The sequence $(G(n,1/2))_{n\in\omega}$ induces $G(\omega,1/2)$ on $\mathrm{Str}_L(\omega)$.

Invariant measures

The space Str_L also comes equipped with a natural action of S_∞, the permutation group of ω.

$\sigma \in S_\infty$ acts on a structure M with domain ω by permuting the domain:

$$\sigma(M) \models R(\overline{a}) \iff M \models R(\sigma^{-1}(\overline{a}))$$

Note:

- If $N = \sigma(M)$, then $\sigma\colon M \to N$ is an isomorphism.
- The orbit of M is $\mathrm{Iso}(M) = \{N \in \mathrm{Str}_L(\omega) \mid M \cong N\}$.

Invariant measures

The space Str_L also comes equipped with a natural action of S_∞, the permutation group of ω.

$\sigma \in S_\infty$ acts on a structure M with domain ω by permuting the domain:

$$\sigma(M) \models R(\overline{a}) \iff M \models R(\sigma^{-1}(\overline{a}))$$

Note:

- If $N = \sigma(M)$, then $\sigma\colon M \to N$ is an isomorphism.
- The orbit of M is $\mathrm{Iso}(M) = \{N \in \mathrm{Str}_L(\omega) \mid M \cong N\}$.

To show that a Borel probability measure on $\mathrm{Str}_L(\omega)$ is invariant for the action of S_∞, it suffices to check:

$$\mu([\varphi(\overline{a})]) = \mu([\varphi(\sigma(\overline{a}))])$$

for all quantifier-free $\varphi(\overline{x})$, $\overline{a} \in \omega$, and $\sigma \in S_\infty$.

Invariant measures

Definition

$(\mu_n)_{n\in\omega}$ is *coherent* if each μ_n is a probability measure on $\mathrm{Str}_L(n)$, and:

1. For all $\varphi(\overline{x})$ quantifier-free, all $\overline{a}$ from $[n]$, and all $n \leq m$, $\mu_n([\varphi(\overline{a})]_n) = \mu_m([\varphi(\overline{a})]_m)$.
2. For all $\varphi(\overline{x})$ quantifier-free, all $\overline{a}$ from $[n]$, and all $\sigma \in S_n$, $\mu_n([\varphi(\overline{a})]_n) = \mu_n([\varphi(\sigma(\overline{a}))]_n)$.
3. For all $\varphi(\overline{x})$ and $\psi(\overline{y})$ quantifier-free and $\overline{a}$ and $\overline{b}$ *disjoint* from $[n]$, $\mu_n([\varphi(\overline{a}) \wedge \psi(\overline{b})]_n) = \mu_n([\varphi(\overline{a})]_n)\mu_n([\psi(\overline{b})]_n)$.

By condition (2), the measure on $\mathrm{Str}_L(\omega)$ induced by a coherent sequence is S_∞-invariant.

Quantifier-free limits of finite structures

Convention: I will refer to S_∞-invariant Borel probability measures on $\mathrm{Str}_L(\omega)$ simply as *invariant measures*.

Definition

- Let A be a finite structure, and let $\varphi(\overline{x})$ be a quantifier-free formula in n free variables. Define $P(\varphi; A) = \frac{|\{\overline{a} \in A^n \mid A \models \varphi(\overline{a})\}|}{|A|^n}$.
- A sequence $(A_n)_{n\in\omega}$ of finite structures with $\lim_{n\to\infty} |A_n| = \infty$ *q.f.-converges* if $\lim_{n\to\infty} P(\varphi; A_n)$ exists for all quantifier-free φ.
- Such a convergent sequence assigns a limiting probability to every quantifier-free formula. There is a unique invariant measure μ on $\mathrm{Str}_L(\omega)$ which encodes these limiting probabilities, and we say $(A_n)_{n\in\omega}$ q.f.-converges to μ.

Example: The Paley graphs $(\mathbb{F}_q^\times, \{(x, y) \mid \exists z, z^2 = x - y\})$ for q a prime power, $q \equiv 1 \pmod 4$, q.f.-converge to $G(\omega, 1/2)$.

This is the kind of convergence captured by graph limits / graphons.

Ergodic structures

Fact

For an invariant measure μ on $\mathrm{Str}_L(\omega)$, the following are equivalent:

1. *There is a sequence of finite L-structures, $(A_n)_{n\in\omega}$ which q.f.-converges to μ.*
2. *For any quantifier-free formulas $\varphi(\overline{x})$ and $\psi(\overline{y})$ and any disjoint tuples $\overline{a}$ and $\overline{b}$ from ω,*

$$\mu([\varphi(\overline{a}) \wedge \psi(\overline{b})]) = \mu([\varphi(\overline{a})])\mu([\psi(\overline{b})])$$

3. *μ is ergodic for the action of S_∞ (i.e. if X is a Borel set such that $\mu(X \triangle \sigma(X)) = 0$ for all $\sigma \in S_\infty$, then $\mu(X) = 0$ or 1).*

Definition (Ackerman–Freer–K.–Patel)

An *ergodic structure* is an invariant measure on $\mathrm{Str}_L(\omega)$ which satisfies the three equivalent conditions in the fact.

Ergodic structures

Definition

$(\mu_n)_{n\in\omega}$ is *coherent* if each μ_n is a probability measure on $\mathrm{Str}_L(n)$, and:

1. For all $\varphi(\overline{x})$ quantifier-free, all $\overline{a}$ from $[n]$, and all $n \leq m$, $\mu_n([\varphi(\overline{a})]_n) = \mu_m([\varphi(\overline{a})]_m)$.
2. For all $\varphi(\overline{x})$ quantifier-free, all $\overline{a}$ from $[n]$, and all $\sigma \in S_n$, $\mu_n([\varphi(\overline{a})]_n) = \mu_n([\varphi(\sigma(\overline{a}))]_n)$.
3. For all $\varphi(\overline{x})$ and $\psi(\overline{y})$ quantifier-free and $\overline{a}$ and $\overline{b}$ *disjoint* from $[n]$, $\mu_n([\varphi(\overline{a}) \wedge \psi(\overline{b})]_n) = \mu_n([\varphi(\overline{a})]_n)\mu_n([\psi(\overline{b})]_n)$.

By condition (3), the invariant measure on $\mathrm{Str}_L(\omega)$ induced by a coherent sequence is an ergodic structure.

More on ergodic structures

Not all ergodic structures concentrate on Fraïssé limits:

Theorem (Ackerman–Freer–Patel)

Let M be a countable structure. The following are equivalent:

1. *M has trivial (group-theoretic) acl: For every finite tuple $\overline{a}$ from M and every $b \in M$, there is an automorphism $\sigma \in \mathrm{Aut}(M)$ such that $\sigma(\overline{a}) = a$ and $\sigma(b) \neq b$.*
2. *There is an invariant measure on $\mathrm{Str}_L(\omega)$ such that $\mu(\mathrm{Iso}(M)) = 1$.*
3. *There is an ergodic structure such that $\mu(\mathrm{Iso}(M)) = 1$.*

More on ergodic structures

Not all ergodic structures concentrate on Fraïssé limits:

Theorem (Ackerman–Freer–Patel)

Let M be a countable structure. The following are equivalent:

1. *M has trivial (group-theoretic) acl: For every finite tuple $\overline{a}$ from M and every $b \in M$, there is an automorphism $\sigma \in \mathrm{Aut}(M)$ such that $\sigma(\overline{a}) = a$ and $\sigma(b) \neq b$.*
2. *There is an invariant measure on $\mathrm{Str}_L(\omega)$ such that $\mu(\mathrm{Iso}(M)) = 1$.*
3. *There is an ergodic structure such that $\mu(\mathrm{Iso}(M)) = 1$.*

Not all ergodic structures give measure 1 to a single isomorphism class:

In joint work with Ackerman, Freer, & Patel (*Properly Ergodic Structures*) we characterized theories T such that there exists an ergodic structure μ with $\mu(\mathrm{Mod}(T)) = 1$ but $\mu(\mathrm{Iso}(M)) = 0$ for all $M \models T$.

Strongly pseudofinite theories

Definition

A theory T is *strongly pseudofinite* if:

1. There is a Fraïssé class K such that $T = T_K$.
2. There is a coherent sequence of measures $(\mu_n)_{n\in\omega}$ which has a zero-one law, and $T^{\text{a.s.}} = T$.

"The generic theory T_K is pseudofinite witnessed by a zero-one law for a reasonable sequence of measures."

Fact (Hill)

It follows from (1) and (2) that if μ is the ergodic structure induced by $(\mu_n)_{n\in\omega}$, then $\mu(\mathrm{Iso}(M_K)) = 1$.

In other words, all three of our limit notions coincide on T.

Example: $\mathrm{Th}(\mathcal{R})$ is strongly pseudofinite.

Strong pseudofiniteness and full amalgamation

Theorem (K.)

If K is a Fraïssé class with full amalgamation, T_K is strongly pseudofinite.

Fact

If T is strongly pseudofinite, and T' is a reduct of T which is also the generic theory of a Fraïssé class, then T' is strongly pseudofinite.

Strong pseudofiniteness and full amalgamation

Theorem (K.)

If K is a Fraïssé class with full amalgamation, T_K is strongly pseudofinite.

Fact

If T is strongly pseudofinite, and T' is a reduct of T which is also the generic theory of a Fraïssé class, then T' is strongly pseudofinite.

Examples:

1. Directed graphs, hypergraphs, directed hypergraphs
2. Bipartite graphs (with the partition named by unary predicates)
3. Simplicial complexes (where n-cells are instances of n-ary relations)
4. 3-hypergraphs in which every tetrahedron has an even number of faces (this is a reduct of the random graph which lacks full amalgamation)

Question

Is every strongly pseudofinite theory a reduct of the generic theory of a Fraïssé class with full amalgamation?

Disjoint n-amalgamation

Notation: $\mathcal{P}^-([n]) = \mathcal{P}([n]) \setminus \{[n]\}$.
We view $\mathcal{P}^-([n])$ and $\mathcal{P}([n])$ as poset categories with a unique arrow $X \to Y$ if and only if $X \subseteq Y$.

Let K be a Fraïssé class, viewed as a category where arrows are embeddings.

A functor F from $\mathcal{P}^-([n])$ or $\mathcal{P}([n])$ to K *preserves disjointness* if for all Z in the domain category of F and all $X, Y \subseteq Z$, the images of $F(X)$ and $F(Y)$ in $F(Z)$ are disjoint over the image of $F(X \cap Y)$ in $F(Z)$.

K has *disjoint n-amalgamation* if every functor $F\colon \mathcal{P}^-([n]) \to K$ which preserves disjointness can be extended to a functor $\widehat{F}\colon \mathcal{P}^-([n]) \to K$ which preserves disjointness.

K has *full amalgamation* if it has disjoint n-amalgamation for all $n \in \omega$.

Disjoint 2-amalgamation

Disjoint 2-amalgamation is often called "strong amalgamation":

$$\begin{tikzcd}
 & A_{\{0,1\}} & \\
A_{\{0\}} \arrow[ur, dashed, "f'"] & & A_{\{1\}} \arrow[ul, dashed, "g'"'] \\
 & A_{\emptyset} \arrow[ul, "f"] \arrow[ur, "g"'] &
\end{tikzcd}$$

...and $f'(A_{\{0\}}) \cap g'(A_{\{1\}}) = f'(f(A_\emptyset)) = g'(g'(A_\emptyset))$ in D.

Fact

A Fraïssé class has disjoint 2-amalgamation if and only if its generic theory T_K has trivial (group-theoretic) definable closure, equivalently trivial (model-theoretic) algebraic closure.

Disjoint 3-amalgamation

$$
\begin{tikzcd}
 & & A_{\{0,1,2\}} & & \\
A_{\{0,1\}} \arrow[urr, dashed] & & A_{\{0,2\}} \arrow[u, dashed] & & A_{\{1,2\}} \arrow[ull, dashed] \\
A_{\{0\}} \arrow[u] \arrow[urr] & & A_{\{1\}} \arrow[ull] \arrow[urr] & & A_{\{2\}} \arrow[ull] \arrow[u] \\
 & & A_{\emptyset} \arrow[ull] \arrow[u] \arrow[urr] & &
\end{tikzcd}
$$

Examples of failure: Let $A_X = \{a_i \mid i \in X\}$:

- $K =$ finite triangle-free graphs. $a_1 E a_2$, $a_2 E a_3$, $a_1 E a_3$.
- $K =$ finite partial orders. $a_1 < a_2$, $a_2 < a_3$, $a_3 < a_1$.
- $K =$ finite equivalence relations. $a_1 E a_2$, $a_2 E a_3$, $\neg a_1 E a_3$.

A random construction

Theorem (K.)

If K is a Fraïssé class with full amalgamation, T_K is strongly pseudofinite.

It suffices to define a coherent sequence of measures $(\mu_n)_{n\in\omega}$ which has a zero-one law with $T^{\mathrm{a.s.}} = T_K$.

I will describe these measures as random constructions of a L-structures with domain $[n]$ for every n, built "from the bottom up".

A random construction

Theorem (K.)

If K is a Fraïssé class with full amalgamation, T_K is strongly pseudofinite.

First, pick the atomic diagram of each element, uniformly at random from those consistent with K.

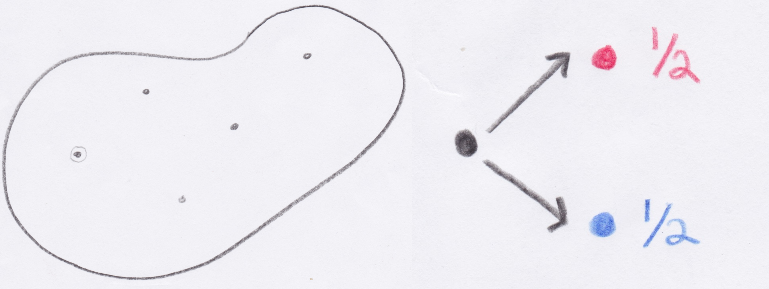

A random construction

Theorem (K.)

If K is a Fraïssé class with full amalgamation, T_K is strongly pseudofinite.

Next, pick the atomic diagram of each pair, uniformly at random from those consistent with K and extending the atomic diagrams assigned to the singletons (disjoint 2-amalgamation implies that the set of choices is non-empty).

A random construction

Theorem (K.)

If K is a Fraïssé class with full amalgamation, T_K is strongly pseudofinite.

Continue in this way, assigning the atomic diagram of each subset of size n uniformly at random from those consistent with K extending the diagrams assigned to the subsets of size $n-1$.

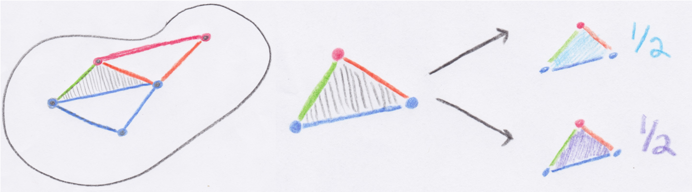

Full amalgamation ensures that we never get stuck, and that all choices could be made as independently as possible.

The zero-one law

It remains to show that the μ_n have a zero-one law with $T^{\text{a.s.}} = T_K$. The proof is a simple generalization of the proof of the zero-one law for finite graphs.

We have described a sequence of measures μ_n on $\mathrm{Str}_L(n)$.

- Since we always build structures in K, it suffices to show that $\lim_{n\to\infty} \mu_n([\varphi]_n) = 1$ when φ is an extension axiom
$$\forall \overline{x}\, \exists y\, (\theta_A(\overline{x}) \to \theta_B(\overline{x}, y)).$$
- For any $\overline{a}$ from $[n]$, if $\theta_A(\overline{a})$, then for any other b, there is a positive probability $\varepsilon > 0$ that $\theta_B(\overline{a}, b)$.
- Conditioned on $[\theta_A(\overline{a})]$, for $b \neq b'$ not in $\overline{a}$, $[\theta_B(\overline{a}, b)]$ and $[\theta_B(\overline{a}, b')]$ are independent. So the conditional probability of $[\neg\exists y\, \theta_B(\overline{a}, b)]$ is $(1-\varepsilon)^{n-|A|}$.
- So $\mu_n([\neg\varphi]_n) \leq n^{|A|}(1-\varepsilon)^{n-|A|} \to 0$ as $n \to \infty$.

The Aldous–Hoover–Kallenberg representation

To understand strongly pseudofinite theories, we need to understand the role of the limiting ergodic structure.

Key tool: a vast generalization of de Finetti's theorem.

The Aldous–Hoover–Kallenberg representation

To understand strongly pseudofinite theories, we need to understand the role of the limiting ergodic structure.
Key tool: a vast generalization of de Finetti's theorem.

Setup:

- $(\xi_A)_{A \in \mathcal{P}_{\mathrm{fin}}(\omega)}$ independent random variables, uniform on $[0,1]$.
- View a non-redundant tuple $a_0, \ldots, a_{n-1} \in \omega$ as an injective function $i \colon [n] \to \omega$.
- Denote by $\widehat{\xi_{\overline{a}}}$ the family of random variables $(\xi_{i[X]})_{X \in \mathcal{P}([n])}$.

Definition

An *AHK system* is a collection of measurable functions

$$(f_n \colon [0,1]^{\mathcal{P}([n])} \to \mathrm{Str}_L(n))_{n \in \omega}$$

satisfying some coherence conditions.

The Aldous–Hoover–Kallenberg representation

Definition

An *AHK system* is a collection of measurable functions

$$(f_n\colon [0,1]^{\mathcal{P}([n])} \to \mathrm{Str}_L(n))_{n\in\omega}$$

satisfying some coherence conditions.

An AHK system allows us to define a structure in $\mathrm{Str}_L(\omega)$ at random, by defining the induced structure on a tuple $\overline{a}$ of length n to be $f_n(\widehat{\xi_{\overline{a}}})$. The coherence conditions ensure that this is well-defined.

The Aldous–Hoover–Kallenberg representation

Definition

An *AHK system* is a collection of measurable functions

$$(f_n\colon [0,1]^{\mathcal{P}([n])} \to \mathrm{Str}_L(n))_{n\in\omega}$$

satisfying some coherence conditions.

An AHK system allows us to define a structure in $\mathrm{Str}_L(\omega)$ at random, by defining the induced structure on a tuple $\overline{a}$ of length n to be $f_n(\widehat{\xi_{\overline{a}}})$. The coherence conditions ensure that this is well-defined.

Definition

If μ is the induced probability measure on $\mathrm{Str}_L(\omega)$, we say $(f_n)_{n\in\omega}$ is an *AHK representation* of μ.

Theorem (Aldous, Hoover, Kallenberg (in different contexts))

Every invariant probability measure μ on Str_L has an AHK representation.

The ergodic case

If $\overline{a}$ and $\overline{b}$ are tuples with intersection $\overline{c}$, then $f_n(\widehat{\xi_{\overline{a}}})$ and $f_m(\widehat{\xi_{\overline{b}}})$ are conditionally independent over $\widehat{\xi_{\overline{c}}}$ ("hidden information at $\overline{c}$").

The ergodic case

If $\overline{a}$ and $\overline{b}$ are tuples with intersection $\overline{c}$, then $f_n(\widehat{\xi_{\overline{a}}})$ and $f_m(\widehat{\xi_{\overline{b}}})$ are conditionally independent over $\widehat{\xi_{\overline{c}}}$ ("hidden information at $\overline{c}$").

Ergodicity corresponds to "no hidden information at $\emptyset$".

Theorem (Aldous, for exchangeable arrays)

Let μ be an invariant measure on $\mathrm{Str}_L(\omega)$. The following are equivalent:

1. *μ is an ergodic structure (recall: ergodic for the action of S_∞).*
2. *μ is "dissociated": For any quantifier-free formulas $\varphi(\overline{x})$ and $\psi(\overline{y})$ and any disjoint tuples $\overline{a}$ and $\overline{b}$ from ω,*

$$\mu([\varphi(\overline{a}) \wedge \psi(\overline{b})]) = \mu([\varphi(\overline{a})])\mu([\psi(\overline{b})]).$$

3. *μ has an AHK representation which does not depend on $\xi_\emptyset$.*

Idea of a proof

Following Austin *On exchangeable random variables and the statistics of large graphs and hypergraphs*, but translated to the setting of L-structures.

Idea of a proof

Following Austin *On exchangeable random variables and the statistics of large graphs and hypergraphs*, but translated to the setting of L-structures.

- Let Ω be a disjoint copy of ω.
- By a bijection $\omega \to \omega \cup \Omega$, transfer μ to a measure $\widehat{\mu}$ on $\mathrm{Str}_L(\omega \cup \Omega)$.

Idea of a proof

Following Austin *On exchangeable random variables and the statistics of large graphs and hypergraphs*, but translated to the setting of L-structures.

- Let Ω be a disjoint copy of ω.
- By a bijection $\omega \to \omega \cup \Omega$, transfer μ to a measure $\widehat{\mu}$ on $\mathrm{Str}_L(\omega \cup \Omega)$.
- Pick a structure $M_\emptyset$ with domain Ω according to $\widehat{\mu}$ ($\xi_\emptyset$).

Idea of a proof

Following Austin *On exchangeable random variables and the statistics of large graphs and hypergraphs*, but translated to the setting of L-structures.

- Let Ω be a disjoint copy of ω.
- By a bijection $\omega \to \omega \cup \Omega$, transfer μ to a measure $\widehat{\mu}$ on $\mathrm{Str}_L(\omega \cup \Omega)$.
- Pick a structure $M_\emptyset$ with domain Ω according to $\widehat{\mu}$ ($\xi_\emptyset$).
- For each $a \in \omega$, pick a structure $M_{\{a\}}$ with domain $\{a\} \cup \Omega$ according to $\widehat{\mu}$, conditionally independently over $M_\emptyset$ ($\xi_{\{a\}}$).

Idea of a proof

Following Austin *On exchangeable random variables and the statistics of large graphs and hypergraphs*, but translated to the setting of L-structures.

- Let Ω be a disjoint copy of ω.
- By a bijection $\omega \to \omega \cup \Omega$, transfer μ to a measure $\widehat{\mu}$ on $\mathrm{Str}_L(\omega \cup \Omega)$.
- Pick a structure $M_\emptyset$ with domain Ω according to $\widehat{\mu}$ ($\xi_\emptyset$).
- For each $a \in \omega$, pick a structure $M_{\{a\}}$ with domain $\{a\} \cup \Omega$ according to $\widehat{\mu}$, conditionally independently over $M_\emptyset$ ($\xi_{\{a\}}$).
- Continue building from the bottom up: For each $A \subseteq_{\text{fin}} \omega$, pick a structure M_A with domain $A \cup \{\Omega\}$ according to $\widehat{\mu}$, conditionally independently over the $(M_B)_{B \subsetneq A}$ for all $A \subseteq_{\text{fin}} \omega$ (ξ_A).

Idea of a proof

Following Austin *On exchangeable random variables and the statistics of large graphs and hypergraphs*, but translated to the setting of L-structures.

- Let Ω be a disjoint copy of ω.
- By a bijection $\omega \to \omega \cup \Omega$, transfer μ to a measure $\widehat{\mu}$ on $\mathrm{Str}_L(\omega \cup \Omega)$.
- Pick a structure $M_\emptyset$ with domain Ω according to $\widehat{\mu}$ ($\xi_\emptyset$).
- For each $a \in \omega$, pick a structure $M_{\{a\}}$ with domain $\{a\} \cup \Omega$ according to $\widehat{\mu}$, conditionally independently over $M_\emptyset$ ($\xi_{\{a\}}$).
- Continue building from the bottom up: For each $A \subseteq_{\text{fin}} \omega$, pick a structure M_A with domain $A \cup \{\Omega\}$ according to $\widehat{\mu}$, conditionally independently over the $(M_B)_{B \subsetneq A}$ for all $A \subseteq_{\text{fin}} \omega$ (ξ_A).
- Show by induction that this process agrees with $\widehat{\mu}$ at every stage.

Idea of a proof

Following Austin *On exchangeable random variables and the statistics of large graphs and hypergraphs*, but translated to the setting of L-structures.

- Let Ω be a disjoint copy of ω.
- By a bijection $\omega \to \omega \cup \Omega$, transfer μ to a measure $\widehat{\mu}$ on $\mathrm{Str}_L(\omega \cup \Omega)$.
- Pick a structure $M_\emptyset$ with domain Ω according to $\widehat{\mu}$ ($\xi_\emptyset$).
- For each $a \in \omega$, pick a structure $M_{\{a\}}$ with domain $\{a\} \cup \Omega$ according to $\widehat{\mu}$, conditionally independently over $M_\emptyset$ ($\xi_{\{a\}}$).
- Continue building from the bottom up: For each $A \subseteq_{\mathrm{fin}} \omega$, pick a structure M_A with domain $A \cup \{\Omega\}$ according to $\widehat{\mu}$, conditionally independently over the $(M_B)_{B \subsetneq A}$ for all $A \subseteq_{\mathrm{fin}} \omega$ (ξ_A).
- Show by induction that this process agrees with $\widehat{\mu}$ at every stage.
- By invariance, the random substructure with domain ω agrees with μ.
- Use standard probability theory tricks to replace all random choices above with random variables on $[0, 1]$.

The case of graphons

Suppose $L = \{E\}$ and μ is an ergodic structure giving measure 1 to the class of graphs. Let $(f_n)_{n\in\omega}$ be an AHK representation for μ.

- The language is binary, so f_n is irrelevant for $n \geq 3$.
- There is only one graph each of size 0 and size 1, so f_0 and f_1 are irrelevant.
- μ is ergodic, so f_2 does not depend on $\xi_\emptyset$.
- For $a, b \in \omega$, $f_2(\xi_{\{a\}}, \xi_{\{b\}}, \xi_{\{a,b\}})$ says either "edge" or "no edge".
- The actual value of $\xi_{\{a,b\}}$ is irrelevant: what matters is probability p, given $\xi_{\{a\}}, \xi_{\{b\}} \in [0,1]$ that $f_2(\xi_{\{a\}}, \xi_{\{b\}}, \xi_{\{a,b\}}) =$ "edge".
- Set $\widehat{f}(\xi_{\{a\}}, \xi_{\{b\}}) = p$.

$\widehat{f}$ is a *graphon*: a (a.s.) symmetric measurable function $[0,1]^2 \to [0,1]$.

AHK representations are the proper generalization of graphons to general relational languages. In specific cases, they can be simplified by an analysis like the above.

MS-Measurability

The rest of this talk is about on-going work with Cameron Hill.

Definition

An AHK system $(f_n)_{n\in\omega}$ is *fully independent* if whenever $\overline{a}$ and $\overline{b}$ are tuples intersecting in $\overline{c}$, then $f_n(\widehat{\xi_{\overline{a}}})$ and $f_m(\widehat{\xi_{\overline{b}}})$ are conditionally independent over $f_k(\widehat{\xi_{\overline{c}}})$.

Slogan: "No hidden information anywhere".

Theorem (Hill-K.)

If T is strongly pseudofinite, and the witnessing ergodic structure μ has a fully independent AHK representation, then T is MS-measurable.

Simplicity

Conjecture (Hill-K.)

If T is strongly pseudofinite, then T has trivial forking:

$$A \underset{C}{\overset{f}{\mathop{\smile\!\!\!\!\!\!|}}} B \iff A \cap B \subseteq C.$$

In particular, it would follow that:

- Every strongly pseudofinite theory is simple of SU-rank 1.
- The generic theory $T_{\mathcal{G}_\triangle}$ of triangle-free graphs is not strongly pseudofinite.

Simplicity

Conjecture (Hill-K.)

If T is strongly pseudofinite, then T has trivial forking:

$$A \underset{C}{\overset{f}{\mathop{\smile\!\!\!\!\!\!\;|}}} B \iff A \cap B \subseteq C.$$

In particular, it would follow that:

- Every strongly pseudofinite theory is simple of SU-rank 1.
- The generic theory $T_{\mathcal{G}_\triangle}$ of triangle-free graphs is not strongly pseudofinite.

Theorem (Hill-K.)

The theory of an equivalence relation with infinitely many infinite classes is not strongly pseudofinite.

Counterexample: equivalence relations

Let M be the equivalence relation with infinitely many infinite classes. Suppose for contradiction that $T = \mathrm{Th}(M)$ is strongly pseudofinite, witnessed by $(\mu_n)_{n\in\omega}$ which cohere to μ.

Counterexample: equivalence relations

Let M be the equivalence relation with infinitely many infinite classes. Suppose for contradiction that $T = \mathrm{Th}(M)$ is strongly pseudofinite, witnessed by $(\mu_n)_{n\in\omega}$ which cohere to μ.

- Any AHK representations of μ essentially has the following form:
  - Fix $(p_i)_{i\in\omega}$ with $0 < p_i < 1$ and $\sum_{i\in\omega} p_i = 1$.
  - The randomness at the level of a singleton $\{a\}$ puts a in an equivalence class C_i with probability p_i.
  - No randomness at the level of pairs (or higher).
  - (In particular, this is a $\{0,1\}$-valued graphon.)

Counterexample: equivalence relations

Let M be the equivalence relation with infinitely many infinite classes. Suppose for contradiction that $T = \mathrm{Th}(M)$ is strongly pseudofinite, witnessed by $(\mu_n)_{n\in\omega}$ which cohere to μ.

- Any AHK representations of μ essentially has the following form:
  - Fix $(p_i)_{i\in\omega}$ with $0 < p_i < 1$ and $\sum_{i\in\omega} p_i = 1$.
  - The randomness at the level of a singleton $\{a\}$ puts a in an equivalence class C_i with probability p_i.
  - No randomness at the level of pairs (or higher).
  - (In particular, this is a $\{0,1\}$-valued graphon.)
- There is some $k > 1$ such that $\lim_{n\to\infty} \mu_n([\forall x\, \exists^{\geq k} y\, xEy]_n) \neq 1$.
  - It suffices to show that there is some $\varepsilon > 0$ such that for any N there is some $n \geq N$ such that $\mu_n([\forall x\, \exists^{\geq k} y\, xEy]_n) < 1 - \varepsilon$.
  - Fix a small ε. Pick p_i small enough so that for some $n \geq N$, $np_i \approx k/2$. This is the expected number of elements in the class C_i.
  - By a Chernoff bound argument, the μ_n-probability that C_i has at least one but less than k elements is at least ε.

(e.g. $k = 10$, $\varepsilon = 1/10$ works)