Lecture Notes: Around Countable Homogeneous Structures

Alex Kruckman

Università di Torino, May/June 2025

Lecture 1: Ages, rich and homogeneous limits

Countable homogeneous structures occupy a beautiful corner of mathematics, which lies at the heart of connections between model theory, combinatorics, descriptive set theory, and permutation group theory.

Here is a brief sketch of how these connections go: Given a class $\mathcal{K}$ of finite (or finitely generated) structures, we study the space of countable $\mathcal{K}$-limits, which are countable structures "built from" the structures in $\mathcal{K}$. Among all $\mathcal{K}$-limits, we identify limits with special model-theoretic properties – in the best case, we find a homogeneous $\mathcal{K}$-limit M, also called the Fraïssé limit of $\mathcal{K}$. Homogeneity implies that the permutation group $\mathrm{Aut}(M)$ is rich and reflects the structure of the original class $\mathcal{K}$.

In these lectures, we cover the following topics:

- Classical Fraïssé theory.
- Strengthenings of amalgamation, and connections to model-theoretic properties.
- Zero-one laws and pseudofiniteness of $\aleph_0$-categorical theories.
- Universal and generic limit structures when no Fraïssé limit exists.

Two excellent general references on Fraïssé theory and homogeneous structures are Macpherson's *A Survey of Homogeneous Structures* and Cameron's *Oligomorphic Permutation Groups*.

We begin by establishing some conventions. "Countable" means finite or countably infinite ($\leq \aleph_0$). $\mathcal{L}$ is always a countable language. "Structure" means $\mathcal{L}$-structure for some fixed countable language $\mathcal{L}$. We allow empty structures.

If A and B are structures, $A \subseteq B$ means that A is a substructure of B. $f\colon A \hookrightarrow B$ means f is an embedding $A \to B$. If $A \subseteq B$, we always have the inclusion embedding $\mathrm{inc}\colon A \hookrightarrow B$. We write $\mathrm{Emb}(A, B)$ for the set of all embeddings $A \hookrightarrow B$.

A structure A is **finitely generated** (f.g.) if there exist $a_1, \ldots, a_n \in A$ such that $A = \langle a_1, \ldots, a_n \rangle$, i.e., A has no proper substructure containing $a_1, \ldots, a_n$.

$A \subseteq_{\text{f.g.}} B$ means that A is a f.g. substructure of B. Since $\mathcal{L}$ is countable, every f.g. structure is countable. If A is f.g., then every embedding $f\colon A \hookrightarrow B$ is determined uniquely by the values $f(a_1), \ldots, f(a_n)$. Thus, when A is f.g. and B is countable, $\text{Emb}(A, B)$ is countable.

Note that if $\mathcal{L}$ is relational (i.e., $\mathcal{L}$ contains only relation symbols), then a structure is f.g. if and only if it is finite.

Definition 1.1. Let M be a countable structure. The **age** of M is

$$\text{Age}(M) = \{A \mid A \text{ is f.g., and } \exists f\colon A \hookrightarrow M\}.$$

The terminology "age" is due to Roland Fraïssé. In Fraïssé's terminology, a structure M is **younger** than a structure N if $\text{Age}(M) \subseteq \text{Age}(N)$.

It follows immediately from the definition that $\text{Age}(M)$ is closed under isomorphism: if $A \in \text{Age}(M)$ and $A \cong B$, then $B \cong \text{Age}(M)$. Since for any embedding $f\colon A \hookrightarrow M$ with A f.g., $A \cong f(A) \subseteq_{\text{f.g.}} M$, we could have equivalently defined $\text{Age}(M)$ as the isomorphism-closure of $\{A \mid A \subseteq_{\text{f.g.}} M\}$.

Definition 1.2. Let $\mathcal{K}$ be any class of f.g. structures. A structure M is a **$\mathcal{K}$-limit** if M is countable and $\text{Age}(M) \subseteq \mathcal{K}$.

Which classes of f.g. structures are ages? The following theorem characterizes these classes. We will take some time proving it in detail, since it illustrates the ideas of the more complex constructions that will occupy us later.

Theorem 1.3. *Let $\mathcal{K}$ be a class of f.g. structures. There exists a countable structure M such that $\mathcal{K} = \text{Age}(M)$ if and only if:*

(1) $\mathcal{K}$ is non-empty and countable up to isomorphism: $1 \leq |\mathcal{K}/{\cong}| \leq \aleph_0$.

(2) $\mathcal{K}$ has the ***hereditary property (HP)****: If $A \hookrightarrow B$, A is f.g., and $B \in \mathcal{K}$, then $A \in \mathcal{K}$.*

(3) $\mathcal{K}$ has the ***joint embedding property (JEP)****: If $A, B \in \mathcal{K}$, then there exists $C \in \mathcal{K}$ and embeddings $f\colon A \hookrightarrow C$ and $g\colon B \hookrightarrow C$.*

Definition 1.4. We say $\mathcal{K}$ is an **age** if it satisfies the condition in Theorem 1.3: $\mathcal{K}$ is non-empty and countable and has HP and JEP.

Note that HP implies that $\mathcal{K}$ is closed under isomorphism: If $A \cong B$, then $A \hookrightarrow B$ and $B \hookrightarrow A$, so $A \in \mathcal{K}$ if and only if $B \in \mathcal{K}$. Many sources define HP in a weaker form: if $A \subseteq_{\text{f.g.}} B$ and $B \in \mathcal{K}$, then $A \in \mathcal{K}$. This property, together with isomorphism-closure, is equivalent to our formulation of HP.

Example 1.5. (1) $\mathcal{K}_{\text{LO}}$ = finite linear orders. $\mathcal{K}_{\text{LO}}$ is an age. In fact, every countably infinite linear order L has $\text{Age}(L) = \mathcal{K}_{\text{LO}}$.

(2) $\mathcal{K}_{\text{G}}$ = finite graphs. $\mathcal{K}_{\text{G}}$ is an age. Every countable graph is a $\mathcal{K}$-limit, but not every countably infinite graph G has $\text{Age}(G) = \mathcal{K}_{\text{G}}$. For example, if G is bipartite, then the triangle graph $\triangle$ is not in $\text{Age}(G)$. Letting R be the random graph, $\text{Age}(R) = \mathcal{K}_{\text{G}}$. Even easier, if we let G be the disjoint union of all finite graphs (up to isomorphism), then $\text{Age}(G) = \mathcal{K}_{\text{G}}$.

(3) $\mathcal{K}_{\text{fields}}$ = finite fields. $\mathcal{K}_{\text{fields}}$ is not an age, because it fails JEP: no two fields of different characteristics embeds in a common field. But for a fixed prime p, $\mathcal{K}_{\text{fields},p}$ = finite fields of characteristic p is an age. $\mathcal{K}_{\text{fields},p} = \text{Age}(\overline{\mathbb{F}_p})$, where $\overline{\mathbb{F}_p}$ is the algebraic closure of the prime field $\mathbb{F}_p$.

(4) For us, a **tree** is a connected acyclic graph. $\mathcal{K}_{\text{trees}}$ = finite trees is not an age, because it fails HP: a substructure of a tree need not be connected. A **forest** is an acyclic graph. Every forest is a disjoint union of trees, namely its connected components. $\mathcal{K}_{\text{forests}}$ = finite forests is an age. Note that $\mathcal{K}_{\text{forests}}$ is the hereditary closure of $\mathcal{K}_{\text{trees}}$.

(5) The condition that $\mathcal{K}$ is countable up to isomorphism is automatic in the case that $\mathcal{L}$ is a finite relational language, since there are only finitely many $\mathcal{L}$-structures of size n up to isomorphism for each $n \in \omega$. However, this condition can fail in general.

   (a) If $\mathcal{L} = \{P_i \mid i \in \omega\}$, with each P_i a unary predicate, then there are $2^{\aleph_0}$-many non-isomorphic $\mathcal{L}$-structures of size 1.

   (b) It is a fact that there are $2^{\aleph_0}$-many f.g. groups up to isomorphism. In fact, there are already $2^{\aleph_0}$-many non-isomorphic groups generated by 2 elements.

   (c) Let $\mathcal{L} = \{P, f\}$, where P is a unary relation symbol and f is a unary function symbol. For each binary string $s \in 2^\omega$, we can make $\mathbb{N}$ into an $\mathcal{L}$-structure N_s by setting $f(n) = n+1$ and $P(n)$ iff $s(n) = 1$. Note that $N_s = \langle 0 \rangle$. This gives $2^{\aleph_0}$-many non-isomorphic $\mathcal{L}$-structures generated by a single element.

Before proving Theorem 1.3, we isolate the following lemma, which will be useful later.

Lemma 1.6. *Assume $\mathcal{K}$ has HP. Then M is a $\mathcal{K}$-limit if and only if there exists a chain of embeddings*

$$A_0 \xrightarrow{f_0} A_1 \xrightarrow{f_1} A_2 \xrightarrow{f_2} \dots$$

with each $A_i \in \mathcal{K}$ such that $M \cong \varinjlim A_i$.

Here $\varinjlim A_i$ is the **direct limit** (or the **directed colimit**, in category-theoretic terminology) of the chain of embeddings. If we assume each f_i is an inclusion $A_i \subseteq A_{i+1}$ (which we can always do, at the cost of replacing each A_i by an isomorphic copy), then this is just the union $\varinjlim A_i = \bigcup_{i \in \omega} A_i$.

Proof. Assume M is a $\mathcal{K}$-limit. Since M is countable, enumerate $M = (m_i)_{i \in \omega}$. Let $A_i = \langle m_0, \dots, m_{i-1} \rangle \subseteq_{\text{f.g.}} M$. Then each $A_i \in \text{Age}(M) \subseteq \mathcal{K}$. We have $A_0 \subseteq A_1 \subseteq A_2 \subseteq \dots$, and $M = \bigcup_{i<\omega} A_i \cong \varinjlim A_i$.

Conversely, assume $M \cong \varinjlim A_i$ for some chain of embeddings with each $A_i \in \mathcal{K}$. Each A_i embeds in M with image A_i'. We have $A_0' \subseteq A_1' \subseteq A_2' \subseteq \dots$, and $M = \bigcup_{i\in\omega} A'_i$. Since $\mathcal{K}$ is closed under isomorphism (by HP), each $A'_i \in \mathcal{K}$. Since each A'_i is f.g., A'_i is countable, so M is countable.

Suppose $B \in \mathrm{Age}(M)$, and let $g\colon B \hookrightarrow M$. let $b_1, \ldots, b_k$ be generators for B. For each $1 \leq j \leq k$, there exists $i \in \omega$ such that $g(b_j) \in A'_i$. Let N be large enough so that $b_j \in A'_N$ for all $1 \leq j \leq k$. Then $g(B) \subseteq A'_N$, and we have $g\colon B \hookrightarrow A'_N$. By HP, $B \in \mathcal{K}$. Thus M is a $\mathcal{K}$-limit. □

Now we return to prove our theorem characterizing ages.

Proof of Theorem 1.3. Assume $\mathcal{K} = \mathrm{Age}(M)$.

(1) M has at least one f.g. substructure, namely $\langle\varnothing\rangle$, so $\mathcal{K}$ is non-empty. For all $A \in \mathrm{Age}(M)$, there exists $f\colon A \hookrightarrow M$, so $A \cong f(A) \subseteq_{\text{f.g.}} M$. But since M is countable, $\{B \mid B \subseteq_{\text{f.g.}} M\}$ is countable, so $K/{\cong}$ is countable.

(2) Suppose $f\colon A \hookrightarrow B$, A is f.g., and $B \in \mathrm{Age}(M)$. Let $g\colon B \hookrightarrow M$. Then $g \circ f\colon A \hookrightarrow M$, so $A \in \mathrm{Age}(M)$. Thus $\mathrm{Age}(M)$ has HP.

(3) Suppose $A, B \in \mathrm{Age}(M)$. Let $f\colon A \hookrightarrow M$ and $g\colon B \hookrightarrow M$. Let $C = \langle f(A) \cup g(B)\rangle \subseteq_{\text{f.g.}} M$. Then $C \in \mathrm{Age}(M)$ and $f(A) \subseteq C$ and $g(B) \subseteq C$, so $f\colon A \hookrightarrow C$ and $g\colon B \hookrightarrow C$. Thus $\mathrm{Age}(M)$ has JEP.

Conversely, suppose $\mathcal{K}$ satisfies the three conditions in the theorem. By (1), let $(A_i)_{i\in\omega}$ be isomorphism representatives for $\mathcal{K}$ (listed with repetitions if $K/{\cong}$ is finite). We build a chain of embeddings in $\mathcal{K}$ by recursion.

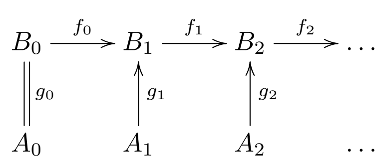

Let $B_0 = A_0$. Given B_i, by JEP there exists some $B_{i+1} \in \mathcal{K}$ with embeddings $f_i\colon B_i \hookrightarrow B_{i+1}$ and $g_{i+1}\colon A_{i+1} \hookrightarrow B_{i+1}$. Let $M = \varinjlim B_i$. By Lemma 1.6 and HP, M is a $\mathcal{K}$-limit, i.e., $\mathrm{Age}(M) \subseteq \mathcal{K}$. Conversely, let $A \in \mathcal{K}$. Then $A \cong A_i$ for some $i < \omega$. We have $A \cong A_i \hookrightarrow B_i \hookrightarrow M$, so $A \in \mathrm{Age}(M)$. Thus $\mathrm{Age}(M) = \mathcal{K}$. □

Among all possible $\mathcal{K}$-limits, we are interested in certain very rich ones (in various senses).

Definition 1.7. Let $\mathcal{K}$ be an age. A $\mathcal{K}$-limit M is:

- **$\mathcal{K}$-rich** if for all $f\colon A \hookrightarrow M$ and $g\colon A \hookrightarrow B$ with $A, B \in \mathcal{K}$, there exists $h\colon B \hookrightarrow M$ such that $h \circ g = f$.

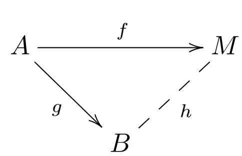

- **$\mathcal{K}$-homogeneous** if $\mathrm{Age}(M) = \mathcal{K}$ and for all $A, A' \subseteq_{\text{f.g.}} M$ and $f\colon A \cong A'$ an isomorphism, there exists $\sigma \in \mathrm{Aut}(M)$ such that $\sigma|_A = f$.
- **$\mathcal{K}$-universal** if every $\mathcal{K}$-limit embeds in M.

If the class $\mathcal{K}$ is clear from context, we drop the prefix and simply write rich, homogeneous, and universal. $\mathcal{K}$-homogeneous structures are sometimes called **ultrahomogeneous**, to distinguish this notion from other notions of homogeneity in model theory.

Lemma 1.8. *If $\mathcal{K}$ is an age and M is a rich $\mathcal{K}$-limit, then* $\mathrm{Age}(M) = \mathcal{K}$.

Proof. Since M is a $\mathcal{K}$-limit, $\mathrm{Age}(M) \subseteq \mathcal{K}$. Conversely, let $A \in \mathcal{K}$, and let $E = \langle \varnothing \rangle \subseteq_{\text{f.g.}} M$. By JEP, there exists some $B \in \mathcal{K}$ and embeddings $f\colon A \hookrightarrow B$ and $g\colon E \hookrightarrow B$. By richness (applied to g), there exists $h\colon B \hookrightarrow M$. Then $h \circ f\colon A \hookrightarrow M$, so $A \in \mathrm{Age}(M)$. □

Closely related to the above proof is the fact that JEP implies that $\mathcal{K}$ contains exactly one structure generated by $\varnothing$ up to isomorphism.

Example 1.9. Note that we include the condition $\mathrm{Age}(M) = \mathcal{K}$ in the definition of $\mathcal{K}$-homogeneous. Recall that $\mathcal{K}_{\mathrm{G}}$ is the class of all finite graphs. Let C_ω be the complete graph on ω. Then C_ω is a $\mathcal{K}_{\mathrm{G}}$-limit, and it is homogeneous (for its age, $\mathrm{Age}(C_\omega) =$ the class of all finite complete graphs). But by our definition, it is not $\mathcal{K}_G$-homogeneous.

Definition 1.10. Let $A \subseteq B$. We say B is a **one-point-extension** of A if there exists $b \in B$ such that $B = \langle A \cup \{b\} \rangle$.

Lemma 1.11. *A $\mathcal{K}$-limit M is rich if and only if for all one-point-extensions $A \subseteq B$ with $A, B \in \mathcal{K}$ and $f\colon A \hookrightarrow M$, there exists $g\colon B \hookrightarrow M$ such that $g|_A = f$.*

Proof. One direction easy, since this condition is a special case of richness. In the other direction, suppose we have $f\colon A \hookrightarrow M$ and $g\colon A \hookrightarrow B$ with $A, B \in \mathcal{K}$. Then $g(A) \subseteq_{\text{f.g.}} B$. Let $b_1, \ldots, b_k$ be generators for B, and for $0 \leq i \leq k$, let $B_i = \langle g(A) \cup \{b_1, \ldots, b_i\} \rangle$. Note that $B_0 = g(A)$, $B_k = B$, and for each B_{i+1}, either $B_{i+1} = B_i$ or B_{i+1} is a one-point extension of B_i.

We build a sequence of embeddings $h_i\colon B_i \hookrightarrow M$ by recursion. Let $h_0 = f \circ g^{-1}\colon g(A) \to M$, and note that $h_0 \circ g = f$. Given $h_i\colon B_i \hookrightarrow M$, by our assumption we can find $h_{i+1}\colon B_{i+1} \hookrightarrow M$ such that $h_{i+1}|_{B_i} = h_i$. Finally, we have $h_k\colon B \hookrightarrow M$, and $h_k|_{g(A)} = h_0$, so $h_k \circ g = f$, as desired. □

Example 1.12. The rational order $(\mathbb{Q}, \leq)$ is $\mathcal{K}_{\mathrm{LO}}$-rich. Indeed, let $A \subseteq B$ be a one-point extension of finite linear orders, and let $f\colon A \hookrightarrow \mathbb{Q}$. Suppose $A = \{a_1 < a_2 < \cdots < a_n\}$, and let b be the unique element of $B \setminus A$. If $n = 0$ (i.e., $A = \varnothing$), we can map b to any element of $\mathbb{Q}$. If $b < a_1$, then we can extend f by $b \mapsto f(a_1) - 1 \in \mathbb{Q}$. If $b > a_n$, then we can extend f by $b \mapsto f(a_n) + 1 \in \mathbb{Q}$. If $a_i < b < a_{i+1}$, then we can extend f by $b \mapsto (f(a_i) + f(a_{i+1}))/2 \in \mathbb{Q}$.

Note that the only facts that we used about $\mathbb{Q}$ are that it is non-empty, dense, and has no greatest or least element. The same argument shows that any countably infinite dense linear order without endpoints is a rich $\mathcal{K}_{\mathrm{LO}}$-limit.

Conversely, any rich $\mathcal{K}_{\mathrm{LO}}$-limit L must be a countably infinite dense linear order without endpoints. For example, for any $a < a'$ in L, we can extend the substructure $A = \{a < a'\}$ to $B = \{a < b < a'\}$. By richness, B embeds in L over A, so there is some $b' \in L$ such that $a < b' < a'$.

Theorem 1.13. *Let M and N be rich $\mathcal{K}$-limits. Suppose $A \subseteq_{\text{f.g.}} M$ and $B \subseteq_{\text{f.g.}} N$ and $f\colon A \cong B$ is an isomorphism. Then there is an isomorphism $\varphi\colon M \cong N$ such that $\varphi|_A = f$.*

Proof. Since M and N are countable, we can enumerate them as $M = (m_i)_{i\in\omega}$ and $N = (n_i)_{i\in\omega}$. We define a sequence of isomorphisms $f_i\colon A_i \cong B_i$ with $A_i \subseteq_{\text{f.g.}} M$ and $B_i \subseteq_{\text{f.g.}} N$ such that for all i, $f \subseteq f_i \subseteq f_{i+1}$, $m_i \in A_{2i+1}$, and $n_i \in B_{2i+2}$.

Set $A_0 = A$, $B_0 = B$, and $f_0 = f$.

At odd stage $2i+1$, given $f_{2i}\colon A_{2i} \cong B_{2i}$, let $A_{2i+1} = \langle A_{2i} \cup \{m_i\}\rangle \subseteq_{\text{f.g.}} M$. Since M is a $\mathcal{K}$-limit, $A_{2i}, A_{2i+1} \in \mathcal{K}$.

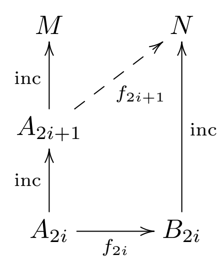

By richness of N, there exists $f_{2i+1}\colon A_{2i+1} \hookrightarrow N$ such that $f_{2i+1}|_{A_{2i}} = f_{2i}$, i.e., $f_{2i} \subseteq f_{2i+1}$. Let $B_{2i+1} = f_{2i+1}(A_{2i+1})$, so $f_{2i+1}\colon A_{2i+1} \cong B_{2i+1}$.

At even stage $2i+2$, the construction is similar, but we go "back", using $f_{2i+1}^{-1}\colon B_{2i+1} \cong A_{2i+1}$ and $B_{2i+2} = \langle B_{2i+1} \cup \{n_i\}\rangle$, and applying richness of M.

Finally, we have $\bigcup_{i\in\omega} A_i = M$, $\bigcup_{i\in\omega} B_i = N$, and $\varphi = \bigcup_{i\in\omega} f_i\colon M \cong N$, with $\varphi|_A = f$. □

The proof of the theorem is a classic "back-and-forth" argument. In fact, the back-and-forth method was originally devised to prove essentially this theorem in the special case of dense linear orders without endpoints. Although Cantor was the first to prove (in 1895) that any two countably infinite dense linear orders without endpoints are isomorphic, and the back-and-forth method is often attributed to him, his proof actually only went "forth" and used a completely different argument to show the constructed embedding was surjective. The first true back-and-forth proofs in the literature seem to be by Huntington (1904) and Hausdorff (1914). The general version of the theorem, for rich $\mathcal{K}$-limits, is due to Fraïssé (1953).

Corollary 1.14. *Let M be a $\mathcal{K}$-limit. Then M is rich if and only if M is homogeneous. Moreover, if M and N are both rich/homogeneous, then $M \cong N$.*

Proof. Suppose M is rich. Then $\mathrm{Age}(M) = \mathcal{K}$ by Lemma 1.8. Let $A, A' \subseteq_{\text{f.g.}} M$ and $f\colon A \cong A'$ an isomorphism. By Theorem 1.13, there is $\sigma \in \mathrm{Aut}(M)$ such that $\sigma|_A = f$. So M is homogeneous.

Conversely, suppose M is homogeneous. Let $f\colon A \hookrightarrow M$ and $g\colon A \hookrightarrow B$ with $A, B \in \mathcal{K}$. Since $\mathrm{Age}(M) = \mathcal{K}$ (this is part of the definition of $\mathcal{K}$-homogeneous), $B \in \mathrm{Age}(M)$, so there exists $h'\colon B \hookrightarrow M$. We may not have $h' \circ g = f$, but there is an isomorphism $\varphi\colon (h' \circ g)(A) \to f(A)$ such that $\varphi \circ h' \circ g = f$. By homogeneity, this isomorphism extends to an automorphism $\sigma \in \mathrm{Aut}(M)$. Let $h = \sigma \circ h'\colon B \hookrightarrow M$. We have $h \circ g = \sigma \circ h' \circ g = \varphi \circ h' \circ g = f$.

Finally, suppose M and N are both rich $\mathcal{K}$-limits. Pick some $A \in \mathcal{K}$. Since $\mathrm{Age}(M) = \mathrm{Age}(N) = \mathcal{K}$ by Lemma 1.8, there are embeddings $f\colon A \hookrightarrow M$ and $g\colon A \hookrightarrow M$, and $f(A) \cong A \cong g(A)$. By Theorem 1.13, $M \cong N$. □

Lecture 2: Fraïssé limits

Definition 2.1. By Corollary 1.14, if $\mathcal{K}$ has a rich or homogeneous limit M, then M is both rich and homogeneous, and it is unique up to isomorphism. We call such a structure M the **Fraïssé limit** of $\mathcal{K}$. We usually denote the Fraïssé limit of $\mathcal{K}$ by $M_{\mathcal{K}}$.

Fraïssé limits also satisfy our third richness property, universality. The proof of this is easier, since we only need to go "forth".

Theorem 2.2. *Suppose $\mathcal{K}$ has a Fraïssé limit $M_{\mathcal{K}}$. Then $M_{\mathcal{K}}$ is $\mathcal{K}$-universal.*

Proof. Let N be a $\mathcal{K}$-limit. By Lemma 1.6, we can write $N = \varinjlim A_i$ for a chain of embedding $f_i\colon A_i \hookrightarrow A_{i+1}$ with each $A_i \in \mathcal{K}$. Since $\mathrm{Age}(M_{\mathcal{K}}) = \mathcal{K}$, there is an embedding $g_0\colon A_0 \hookrightarrow M$. Given an embedding $g_i\colon A_i \hookrightarrow M_{\mathcal{K}}$, by richness there is an embedding $g_{i+1}\colon A_{i+1} \hookrightarrow M_{\mathcal{K}}$ such that $g_{i+1} \circ f_i = g_i$.

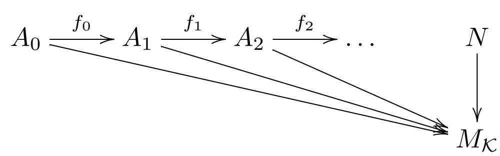

Since $N = \varinjlim A_i$, these coherent embeddings $(g_i)_{i\in\omega}$ induce an embedding $g\colon N \hookrightarrow M_{\mathcal{K}}$. □

Example 2.3. The converse to Theorem 2.2 does not hold. For example, the linear order $\mathbb{Q} \cup \{\infty\}$ is universal, since every countable linear order already embeds in $\mathbb{Q}$. However, it is not homogeneous, since although the substructures $\{0\}$ and $\{\infty\}$ are isomorphic, there is no automorphism moving the greatest element ∞ to 0.

More generally, if $N \subseteq M$ are $\mathcal{K}$-limits and N is universal, then M is universal. But if N is a Fraïssé limit, M is typically not a Fraïssé limit.

Example 2.4. Less trivially, there are classes $\mathcal{K}$ which admit a universal $\mathcal{K}$-limit but no Fraïssé limit. Take, for example, the class $\mathcal{K}_{\text{forests}}$ of finite forests, defined in Example 1.5(4). It is easy to see that every countable tree embeds in the complete countably branching tree $\mathcal{T}$. Since every countable forest is a disjoint union of countably many countable trees, every countable forest embeds in the forest $\mathcal{F}$ formed as the disjoint union of countably many copies of $\mathcal{T}$, and $\mathcal{F}$ is universal for $\mathcal{K}_{\text{forests}}$. Note, however, that if we add a new vertex to $\mathcal{F}$ which is connected to the root of each copy of $\mathcal{T}$, we obtain a tree isomorphic to $\mathcal{T}$. Thus $\mathcal{T}$ is also universal for $\mathcal{K}_{\text{forests}}$.

On the other hand, $\mathcal{K}_{\text{forests}}$ has no Fraïssé limit. Indeed, suppose for contradiction that M were a Fraïssé limit. The path of length 3 embeds in M: $\bullet - \bullet - \bullet - \bullet$. Let A be the substructure of M containing the endpoints of this path: $A = \{a, a'\}$ in $a - \bullet - \bullet - a'$. Note that A is a structure consisting of two disconnected vertices. Now consider the extension $A \hookrightarrow B$ where B is $a - \bullet - a'$. Since B embeds in M over A, M contains both a path of length 3 and a path of length 2 from a to a'. But this creates a cycle of length 5 in M, contradicting the fact that M is acyclic.

Theorem 2.5. *Let $\mathcal{K}$ be an age. Then $\mathcal{K}$ has a Fraïssé limit if and only if $\mathcal{K}$ has the **amalgamation property (AP)**: For all $f_1\colon A \hookrightarrow B_1$ and $f_2\colon A \hookrightarrow B_2$, there exists $C \in \mathcal{K}$, $g_1\colon B_1 \hookrightarrow C$, and $g_2\colon B_2 \hookrightarrow C$ such that $g_1 \circ f_1 = g_2 \circ f_2$.*

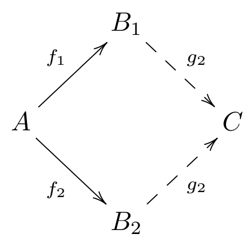

Proof. Suppose first that $\mathcal{K}$ has a Fraïssé limit $M_{\mathcal{K}}$. For AP, let $f_1\colon A \to B_1$ and $f_2\colon A \to B_2$ with $A, B_1, B_2 \in \mathcal{K}$. Since $B_1 \in \text{Age}(M_{\mathcal{K}})$, there is $g_1\colon B_1 \hookrightarrow M_{\mathcal{K}}$. Then we have $(g_1 \circ f_1)\colon A \hookrightarrow M_{\mathcal{K}}$ and $f_2\colon A \hookrightarrow B_2$. By richness, there exists $g_2\colon B_2 \hookrightarrow M_{\mathcal{K}}$ such that $g_2 \circ f_1 = g_1 \circ f_1$. Define $C = \langle g_1(B_1) \cup g_2(B_2)\rangle \subseteq_{\text{f.g.}} M_{\mathcal{K}}$, so $C \in \mathcal{K}$ and we have $g_1\colon B_1 \hookrightarrow C$ and $g_2\colon B_2 \hookrightarrow C$.

Conversely, suppose $\mathcal{K}$ has AP. Let $(A_k)_{k \in \omega}$ be an enumeration of the structures in $\mathcal{K}$, up to isomorphism. We modify the proof of Theorem 1.3 to build a rich $\mathcal{K}$-limit, building a chain of embeddings $f_i\colon B_i \hookrightarrow B_{i+1}$ with each $B_i \in \mathcal{K}$ and defining $M = \varinjlim B_i$. Note that if $i \leq j$, there is an embedding $f_{i,j}\colon B_i \hookrightarrow B_j$, obtained by composing along the chain.

A **task** is an embedding $g\colon C \to A_k$, where A_k is one of our isomorphism representatives for $\mathcal{K}$ and $C \subseteq_{\text{f.g.}} B_i$ for some $i < \omega$.

We will list the tasks by pairs in $\omega \times \omega$. We define tasks (i, j) for all $j \in \omega$ at the end of stage i of the construction, as follows: Having defined B_i, consider all the tasks $g\colon C \to A_k$ with $C \subseteq_{\text{f.g.}} B_i$ and $k \in \omega$. Since B_i is countable, there are countably many $C \subseteq_{\text{f.g.}} B_i$, there are countably many A_k, and $\text{Emb}(C, A_k)$

is countable. Thus there are countably many such g, and we can enumerate them as $(g_j)_{j\in\omega}$. Let task (i,j) be g_j.

Fix a bijection $t\colon \omega \to \omega \times \omega$ such that $t(n) = (i,j)$ with $i \leq n$ for all n (this ensures that task $t(n)$ has already been defined at stage $n+1$). The usual "diagonals" bijection, which enumerates $\omega\times\omega$ as $(0,0)$, $(0,1)$, $(1,0)$, $(0,2)$, $(1,1)$, $(2,0)$,… works. We will complete task $t(n)$ at stage $n+1$ of the construction.

At stage 0, pick any $B_0 \in \mathcal{K}$. Then we define tasks $(0,j)$ for $j \in \omega$, as described above.

At stage $n+1$, we are given B_n, and we seek to construct B_{n+1} and complete task $t(n) = (i,j)$ with $i \leq j$. This task is an embedding $g\colon C \to A_k$, where $C \subseteq_{\text{f.g.}} B_i$. By composing the inclusion $C \hookrightarrow B_i$ with $f_{i,n}\colon B_i \hookrightarrow B_n$, we obtain $f\colon C \hookrightarrow B_n$. Applying AP to f and g, we obtain B_{n+1} and embeddings $g'\colon A \hookrightarrow B_{n+1}$ and $f_n\colon B_n \hookrightarrow B_{n+1}$:

$$\begin{array}{ccccc} C & \xrightarrow{\text{inc}} & B_i & \xrightarrow{f_{i,n}} & B_n \\ {\scriptstyle g}\downarrow & & & & \downarrow{\scriptstyle f_n} \\ A & & \xrightarrow{\quad g' \quad} & & B_{n+1} \end{array}$$

Finally, we define tasks $(n+1,j)$ for $j \in \omega$, as described above.

At the end of this recursive process, we have defined the entire sequence $(B_n)_{n\in\omega}$ and completed every task $t(n)$. Let $M = \varinjlim B_n$. By Lemma 1.6, M is a $\mathcal{K}$-limit. It remains to check richness.

So suppose $f\colon C \to M$ and $g\colon C \to A$ are embeddings, with $A, C \in \mathcal{K}$. Since C is f.g., f factors through $f'\colon C \hookrightarrow B_i$ for some $i < \omega$. After composing with isomorphisms, we can assume that $C \subseteq_{\text{f.g.}} B_i$ and A is A_k for some $k \in \omega$. So $g\colon C \to A$ was added as a task $(i,j) = t(n)$ for some j and n, and at stage $n+1$, we ensured that there was an embedding $g'\colon A \hookrightarrow B_{n+1}$ such that $g' \circ g = f_{i,n+1}|_C$. Composing with the canonical embedding $B_{n+1} \to M$ finishes the proof. □

Definition 2.6. We say $\mathcal{K}$ is a **Fraïssé class** if it is an age which satisfies the condition of Theorem 2.5. Equivalently, $\mathcal{K}$ is non-empty, countable up to isomorphism, and has HP, JEP, and AP.

Example 2.7. Here are some examples of Fraïssé classes and their Fraïssé limits.

- The class of finite sets (in the empty language). The Fraïssé limit is the countably infinite set ω.
- The class of finite linear orders. The Fraïssé limit is the countably infinite dense linear order without endpoints $(\mathbb{Q}, \leq)$.
- The class of finite graphs. The Fraïssé limit is the **random graph** (or the Rado graph) R. It is characterized by the following extension property: For all finite $A, B \subseteq R$ with $A \cap B = \varnothing$, there exists $v \in R$ such that vEa for all $a \in A$ and $\neg vEb$ for all $b \in B$.

- The class of finite triangle-free graphs. The Fraïssé limit is called the **Henson graph** H. It is characterized by the following extension property: For all finite $A, B \subseteq R$ with $A \cap B = \varnothing$ and such that A is an independent set (no edges between vertices in A), there exists $v \in H$ such that vEa for all $a \in A$ and $\neg vEb$ for all $b \in B$.

- The class of finite equivalence relations. The Fraïssé limit is the equivalence relation with countably infinitely many countably infinite classes.

- The class of finite fields of characteristic p. The Fraïssé limit is $\overline{\mathbb{F}_p}$, the algebraic closure of the prime field $\mathbb{F}_p$.

- The class of non-trivial finite Boolean algebras. The Fraïssé limit is the countable atomless Boolean algebra.

Example 2.8. Returning to Example 2.4, the class $\mathcal{K}_{\text{forests}}$ fails to have the amalgamation property. Our argument that there is no homogeneous limit can be translated to a failure of AP. Let $A = \{a, a'\}$ consist of two disconnected vertices. Embed A in B_1: $a - \bullet - a'$ and in B_2: $a - \bullet - \bullet - a'$. Then there is no C amalgamating B_1 and B_2 over A, since any such C must contain a cycle of length 5.

Lecture 3: The finitary case, DAP

To bring Fraïssé limits into the context of first-order model theory, it is convenient to restrict attention to classes where we can define structures and embeddings by first-order formulas.

For $n \in \omega$, an n**-generated structure** is $(A, a_1, \ldots, a_n)$, where $a_1, \ldots, a_n$ are specified generators for A. Two n-generated structures $(A, a_1, \ldots, a_n)$ and $(B, b_1, \ldots, b_n)$ are isomorphic if there exists $f\colon A \cong B$ such that $f(a_i) = b_i$ for all $1 \leq i \leq n$. Equivalently, A and B are isomorphic in the language where the generators are named by n new constant symbols.

Definition 3.1. Let $\mathcal{K}$ be a class of f.g. structures. We say $\mathcal{K}$ is **finitary** if for every n, there are only finitely many n-generated structures in $\mathcal{K}$, up to isomorphism of n-generated structures.

Note that if $\mathcal{L}$ is a finite relational language, then every class of f.g. structures is finitary, since an n-generated structure has size $\leq n$, and there are only finitely many structures of size n up to isomorphism for all n. Note also that if $\mathcal{K}$ is finitary, then K is countable up to isomorphism, which is one of the conditions characterizing ages in Theorem 1.3.

Example 3.2. A class $\mathcal{K}$ of finite structures need not be finitary. For example, the class $\mathcal{K}_{\text{fields},p}$ of finite fields of characteristic p is not finitary. By the Primitive Element Theorem (or the fact that the multiplicative group of a finite field is cyclic), every finite field $\mathbb{F}_{p^n}$ can be written as $\mathbb{F}_p[\alpha] = \langle \alpha \rangle$ for some α. Thus $\mathcal{K}_{\text{fields},p}$ contains infinitely many 1-generated structures up to isomorphism.

Example 3.3. The class of finite Boolean algebras is finitary. Indeed, an n-generated Boolean algebra has size at most 2^{2^n}, and there are only finitely many Boolean algebras of size $\leq 2^{2^n}$ up to isomorphism, since the language is finite.

Example 3.4. Here is an example of a finitary Fraïssé class in an infinite relational language. Let $\mathcal{L} = \{R_n \mid 1 \leq n \in \omega\}$, where each R_n is an n-ary relation symbol. A **simplicial complex** is an $\mathcal{L}$-structure satisfying the following properties:

(1) If $R_n(a_1, \ldots, a_n)$, then $a_i \neq a_j$ for all $i \neq j$.

(2) If $R_n(a_1, \ldots, a_n)$ and $\sigma \in S_n$, then $R_n(a_{\sigma(1)}, \ldots, a_{\sigma(n)})$.

(3) If $R_n(a_1, \ldots, a_n)$, then for all $k < n$, $R_k(a_1, \ldots, a_k)$.

The call $\mathcal{K}_\Delta$ of all simplicial complexes is a finitary Fraïssé class, whose Fraïssé limit is called the **random simplicial complex**. To see that it is finitary, note that if $A \in \mathcal{K}_\Delta$ with $|A| \leq n$, then no relation R_m with $m > n$ can hold of any tuple from A. So to determine A up to isomorphism, we only need to look at a finite sublanguage of $\mathcal{L}$.

Given an n-generated structure $(A, \overline{a})$, we define

$$\mathrm{Diag}(A, \overline{a}) = \{\varphi(\overline{x}) \mid \varphi \text{ is atomic or negated atomic, and } A \models \varphi(\overline{a})\}.$$

Note that $(A, \overline{a}) \cong (B, \overline{b})$ if and only if $\mathrm{Diag}(A, \overline{a}) = \mathrm{Diag}(B, \overline{b})$.

Proposition 3.5. *Suppose $\mathcal{K}$ is a finitary age. Then there is a universal theory $T_\mathcal{K}$ such that for all countable M, $M \models T_\mathcal{K}$ if and only if M is a $\mathcal{K}$-limit.*

Proof. Let $(B, \overline{b})$ be an n-generated structure which is not in $\mathcal{K}$. Since $\mathcal{K}$ is finitary, we can let $(A_1, \overline{a}^1), \ldots, (A_k, \overline{a}^k)$ enumerate the n-generated structures in $\mathcal{K}$ up to isomorphism. Then for all $1 \leq i \leq n$, $(B, \overline{b}) \ncong (A_i, \overline{a}^i)$, so there is some $\chi_i(\overline{x}) \in \mathrm{Diag}(B, \overline{b})$ with $\neg\chi_i(\overline{x}) \in \mathrm{Diag}(A_i, \overline{a}^i)$. Let $\chi_B(\overline{x})$ be $\bigwedge_{i=1}^k \chi_i(\overline{x})$. Note that $B \models \chi_B(\overline{b})$, but for all $1 \leq i \leq k$, $A_i \models \neg\chi_B(\overline{a}^i)$.

Now define

$$T_\mathcal{K} = \{\forall \overline{x} \neg \chi_B(\overline{x}) \mid (B, \overline{b}) \notin \mathcal{K}\}.$$

Suppose M is a $\mathcal{K}$-limit. For any n-generated $(B, \overline{b}) \notin \mathcal{K}$, we show $M \models \forall \overline{x} \neg \chi_B(\overline{x})$. Let $\overline{a} \in M$, and let $A = \langle \overline{a} \rangle$. Then since $A \in \mathcal{K}$, there is some $1 \leq i \leq k$ such that $(A, \overline{a}) \cong (A_i, \overline{a}^i)$, and since $A_i \models \neg\chi_B(\overline{a}^i)$ and χ_B is quantifier-free, $M \models \neg\chi_B(\overline{a})$.

Conversely, suppose M is not a $\mathcal{K}$-limit. Then there is $B \subseteq_{\text{f.g.}} M$ with $B \notin \mathcal{K}$. Let $\overline{b}$ be generators for B. Then $B \models \chi_B(\overline{b})$, and since χ_B is quantifier-free, $M \models \chi_B(\overline{b})$, so $M \not\models \forall \overline{x} \neg\chi_B(\overline{x})$, and $M \not\models T_\mathcal{K}$. □

Proposition 3.6. *Suppose $\mathcal{K}$ is a finitary age, and let $T_\mathcal{K}$ be the universal theory from Proposition 3.5. For every n-generated structure $(A, \overline{a})$ in $\mathcal{K}$, there is a quantifier-free formula $\theta_A(\overline{x})$ such that for all $M \models T_\mathcal{K}$ and $\overline{b}$ from M with $B = \langle \overline{b} \rangle$, $M \models \theta_A(\overline{b})$ if and only if $(B, \overline{b}) \cong (A, \overline{a})$.*

Proof. The proof is similar to that of Proposition 3.5. Let $(A_1, \overline{a}^1), \ldots, (A_k, \overline{a}^k)$ enumerate the n-generated structures in $\mathcal{K}$ which are *not* isomorphic to $(A, \overline{a})$. Let $\theta_A(\overline{x})$ be the conjunction of formulas separating $\mathrm{Diag}(A, \overline{a})$ from $\mathrm{Diag}(A_i, \overline{a}^i)$ for each i, so $A \models \theta_A(\overline{a})$ but $A_i \models \neg\theta_A(\overline{a}^i)$ for each i.

Now given $M \models T_{\mathcal{K}}$ and $\overline{b}$ from M with $B = \langle \overline{b} \rangle$, $M \models \theta_A(\overline{b})$ iff $B \models \theta_A(\overline{b})$ iff $(B, \overline{b}) \cong (A, \overline{a})$. Indeed, since $M \models T_{\mathcal{K}}$, $B \in \mathcal{K}$, so $(B, \overline{b})$ is either isomorphic to $(A, \overline{a})$ or one of the $(A_i, \overline{a}^i)$. □

Given a finitary Fraïssé class $\mathcal{K}$, let $A \subseteq B$ be a one-point extension in $\mathcal{K}$. Choose generators $\overline{a}$ for A and $b \in B$ such that B is generated by $\overline{a}b$. Let $\varphi_{A,B}$ be the sentence

$$\forall \overline{x}(\theta_A(\overline{x}) \to \exists y\, \theta_B(\overline{x}, y)).$$

The sentence $\varphi_{A,B}$ is called a **one-point extension axiom**. Let $T^*_{\mathcal{K}}$ be the theory $T_{\mathcal{K}}$ together with all one-point extension axioms $\varphi_{A,B}$.

Theorem 3.7. *Let $\mathcal{K}$ be a finitary Fraïssé class with no finite upper bound on the size of structures in $\mathcal{K}$. $T^*_{\mathcal{K}}$ is complete and $\aleph_0$-categorical and eliminates quantifiers. Its unique countable model (up to isomorphism) is the Fraïssé limit of $\mathcal{K}$.*

Proof. Let M be a countable model of $T^*_{\mathcal{K}}$. Since $M \models T_{\mathcal{K}}$, M is a $\mathcal{K}$-limit, and the axioms $\varphi_{A,B}$ express that M is rich (using Lemma 1.11). So M is the Fraïssé limit of $\mathcal{K}$. Since there is no finite upper bound on the size of structures in $\mathcal{K}$, M is infinite. By uniqueness of the Fraïssé limit (Corollary 1.14), any two countably infinite models of $T^*_{\mathcal{K}}$ are isomorphic, so $T^*_{\mathcal{K}}$ is $\aleph_0$-categorical.

We have seen that $T^*_{\mathcal{K}}$ is $\aleph_0$-categorical and has no finite models. By Vaught's test, $T^*_{\mathcal{K}}$ is complete.

For quantifier-elimination, let $\psi(\overline{x})$ be any formula, where $\overline{x}$ is a tuple of length n. For any n-tuple $\overline{a}$ from M, let $A = \langle \overline{a} \rangle$. Now let $\overline{a}'$ be any tuple such that $M \models \theta_A(\overline{a}')$. Letting $A' = \langle \overline{a}' \rangle$, we have $(A, \overline{a}) \cong (A', \overline{a}')$, so by homogeneity, there is an automorphism $\sigma \in \mathrm{Aut}(M)$ such that $\sigma(\overline{a}) = \overline{a}'$. It follows that $M \models \psi(\overline{a})$ if and only if $M \models \psi(\overline{a}')$. Thus

$$M \models \forall \overline{x} \left(\psi(\overline{x}) \leftrightarrow \bigvee \theta_{A_i}(\overline{x}) \right)$$

where the disjunction is taken over those n-generated substructures $(A_i, \overline{a}^i)$ such that $M \models \psi(\overline{a}^i)$. Since $\mathcal{K}$ is finitary, there are only finitely many such formulas θ_{A_i}, so $\psi(\overline{x})$ is equivalent to a finite disjunction of quantifier-free formulas.

We have shown that $\psi(\overline{x})$ is equivalent to a quantifier-free formula in M. Since $T^*_{\mathcal{K}}$ is complete, this equivalence holds in all models of $T^*_{\mathcal{K}}$. □

Theorem 3.7 has a converse.

Proposition 3.8. *Suppose T is an $\aleph_0$-categorical theory which eliminates quantifiers. Let $M \models T$ be the unique countably infinite model, and let $\mathcal{K} = \mathrm{Age}(M)$. Then $\mathcal{K}$ is a Fraïssé class with Fraïssé limit M.*

Proof. Since $\mathrm{Age}(M) = \mathcal{K}$, to show that M is the Fraïssé limit of $\mathcal{K}$ (and hence $\mathcal{K}$ is a Fraïssé class), it suffices to show that M is $\mathcal{K}$-homogeneous.

Let $A, B \subseteq_{\text{f.g.}} M$, and assume we have an isomorphism $f\colon A \cong B$. Let $\overline{a}$ be generators for A, and let $\overline{b} = f(\overline{a})$. Then $B = \langle \overline{b} \rangle$, and f is an isomorphism of n-generated structures $(A, \overline{a}) \cong (B, \overline{b})$. We have $\mathrm{Diag}(A, \overline{a}) = \mathrm{Diag}(B, \overline{b})$, so $\mathrm{qftp}(\overline{a}) = \mathrm{qftp}(\overline{b})$. By quantifier elimination, $\mathrm{tp}(\overline{a}) = \mathrm{tp}(\overline{b})$.

Now the unique countably infinite model of an $\aleph_0$-saturated theory is homogeneous in the sense of first-order logic (e.g., since it is both prime and saturated). So there is an automorphism $\sigma \in \mathrm{Aut}(M)$ such that $\sigma(\overline{a}) = \overline{b}$. Since σ agrees with f on the generators $\overline{a}$, $\sigma|_A = f$, as desired. □

Every theory T has an expansion by definitions T' which eliminates quantifiers, by the process of **Morleyization**. Given an $\mathcal{L}$-theory T, we define a language $\mathcal{L}' \supseteq \mathcal{L}$ and an $\mathcal{L}'$-theory T':

$$\mathcal{L}' = \mathcal{L} \cup \{R_{\varphi(\overline{x})} \mid \varphi(\overline{x}) \text{ an } \mathcal{L}\text{-formula}\}$$
$$T' = T \cup \{\forall \overline{x}\, (R_{\varphi(\overline{x})}(\overline{x}) \leftrightarrow \varphi(\overline{x})\}.$$

here $R_{\varphi(\overline{x})}$ is a relation symbol of arity the length of $\overline{x}$. Since T' is an expansion of T by definitions, every model of T has a unique expansion to a model of T', further, every $\mathcal{L}'$-formula is equivalent modulo T' to an $\mathcal{L}$-formula, and every $\mathcal{L}$-formula $\varphi(\overline{x})$ is equivalent modulo T' to the quantifier-free formula $R_{\varphi(\overline{x})}(\overline{x})$. So T' eliminates quantifiers.

If T is $\aleph_0$-categorical, then T' is also $\aleph_0$-categorical. If M is the unique countably infinite model of T, then it has a unique expansion $M' \models T'$, and by Proposition 3.8, M' is the Fraïssé limit of its age (in the language $\mathcal{L}'$). Thus, we have shown that every countable model of an $\aleph_0$-categorical theory is a reduct of a Fraïssé limit.

We turn now to consider certain strengthenings of AP which influence the properties of the Fraïssé limit.

Definition 3.9. A Fraïssé class $\mathcal{K}$ has **disjoint amalgamation (DAP)** if for all $f_1\colon A \hookrightarrow B_1$ and $f_2\colon A \hookrightarrow B_2$ in $\mathcal{K}$, there exists $C \in \mathcal{K}$ and $g_1\colon B_1 \hookrightarrow C$ and $g_2\colon B_2 \hookrightarrow C$ such that $g_1 \circ f_1 = g_2 \circ f_2$, and furthermore, letting $A' = g_1(f_1(A)) = g_2(f_2(A)) \subseteq C$ we have $g_1(B_1) \cap g_2(B_2) = A'$.

Let M be a structure. For $A \subseteq M$ and $b \in M$,

$$\mathrm{Aut}(M/A) = \{\sigma \in \mathrm{Aut}(M) \mid \forall a \in A, \sigma(a) = a\}.$$
$$\mathrm{orb}_A(b) = \{\sigma(b) \mid \sigma \in \mathrm{Aut}(M/B)\}$$
$$\mathrm{ACL}(A) = \{b \in M \mid \mathrm{orb}_A(b) \text{ is finite}\}.$$

The set $\mathrm{ACL}(A)$ is called the **group-theoretic algebraic closure** of A, to distinguish it from the **model-theoretic algebraic closure**:

$$\mathrm{acl}(A) = \{b \in M \mid \exists \varphi(x) \in \mathcal{L}(A) \text{ such that } b \in \varphi(M) \text{ and } \varphi(M) \text{ is finite }\}.$$

We always have $\mathrm{acl}(A) \subseteq \mathrm{ACL}(A)$, since if $b \in \varphi(M)$ and $\varphi(M)$ is finite, $\mathrm{orb}_A(b) \subseteq \varphi(M)$, so $\mathrm{orb}_A(b)$ is finite. When M is the countable model of an $\aleph_0$-categorical theory and $A \subseteq M$ is finite, $\mathrm{acl}(A) = \mathrm{ACL}(A)$.

Theorem 3.10. *Let $\mathcal{K}$ be a Fraïssé class consisting of* finite *structures. Let $M_{\mathcal{K}}$ be the Fraïssé limit of $\mathcal{K}$. Then $\mathcal{K}$ has DAP if and only if* $\mathrm{ACL}(A) = \langle A \rangle$ *for all finite $A \subseteq M_{\mathcal{K}}$.*

To prove the theorem, we will use the following group-theoretic lemma, which comes up surprisingly often in model theory.

Lemma 3.11 (B. H. Neumann's Lemma)**.** *Let G be a group. Assume we can cover G by finitely many cosets of subgroups: $G = \bigcup_{i<n} g_i H_i$, where each $g_i \in G$ and $H_i \leq G$. Then at least one of the groups H_i has finite index in G.*

Proof. We may assume that no proper subset of $\{g_0H_0, \ldots, g_{n-1}H_{n-1}\}$ covers G. Under this assumption, let $H = \bigcap_{i<n} H_i$. We show that $[G : H] \leq n!$, and hence $[G : H_i]$ is finite for all $i < n$.

Let $[n] = \{0, \ldots, n-1\}$. For $X \subseteq [n]$, let $H_X = \bigcap_{i \in X} H_i$. We prove the following claim by induction on m: For all $m \leq n$, if $|X| = n - m$, then $[H_X : H] \leq m!$.

When $m = 0$, $|X| = n$, so $H_X = H$, and $[H_X : H] = 1 \leq 0!$.

Now given $0 < m \leq n$, let $X \subseteq [n]$ with $|X| = n - m$. By minimality, $\{g_iH_i \mid i \in X\}$ does not cover G, so there is some $a \in G$ with $a \notin g_iH_i$ for all $i \in X$. Then for all $i \in X$, aH_i and g_iH_i are disjoint, so g_iH_i is disjoint from aH_X, and $a^{-1}g_iH_i$ is disjoint from H_X. (When $n = m$, $X = \varnothing$, and the conditions on a are vacuous, so we can choose $a = e$.)

Since $\{g_0H_0, \ldots, g_{n-1}H_{n-1}\}$ covers G, so does $\{a^{-1}g_0H_0, \ldots, a^{-1}g_{n-1}H_{n-1}\}$. Thus we can write

$$H_X = \bigcup_{i<n} (H_X \cap a^{-1}g_iH_i).$$

For each $i < n$, $H_X \cap a^{-1}g_iH_i$ is either empty or a coset of $H_X \cap H_i = H_{X\cup\{i\}}$, and we have shown it is empty whenever $i \in X$. Thus, if we remove the empty terms from the union, we are left with at most $n - (n - m) = m$ terms.

When $i \notin X$, $|X \cup \{i\}| = n-(m-1)$, so by induction, $[H_{X\cup\{i\}} : H] \leq (m-1)!$. We have shown that H_X can be covered by at most m cosets of the form $b_iH_{X\cup\{i\}}$, and each of these can be covered by at most $(m-1)!$ cosets of H. So H_X can be covered by at most $m!$ cosets of H, and thus $[H_X : H] \leq m!$.

Finally, when $m = n$, we have $H_X = G$, and the claim tells us $[G : H] \leq n!$, as desired. □

The next result, P. M. Neumann's Lemma, is an immediate consequence of B. H. Neumann's Lemma. One can alternatively deduce B. H. Neumann's Lemma as an easy consequence of P. M. Neumann's Lemma. As an interesting historical note, B. H. Neumann was P. M. Neumann's father.

Lemma 3.12 (P. M. Neumann's Lemma)**.** *Let G be a group acting on a set X. Suppose B and C are finite subsets of X such that each element of B has infinite orbit under the action of G. Then there is some $g \in G$ such that $gB \cap C = \varnothing$.*

Proof. Enumerate B as $b_1, \ldots, b_m$ and C as $c_1, \ldots, c_n$. For each $1 \leq i \leq m$, let $H_i = \mathrm{Stab}(b_i) \leq G$. Since the orbit of b_i is infinite, H_i has infinite index in G.

For each $1 \le i \le m$ and $1 \le j \le n$, $\{g \in G \mid gb_i = c_j\}$ is either empty or a coset $g_{i,j}H_i$ of H_i. Since each H_i has infinite index in G, by Lemma 3.11, the cosets $g_{i,j}H_i$ fail to cover G. Thus we can find some $g \in G$ such that for all i and j, $gb_i \neq c_j$, i.e., $gB \cap C = \varnothing$. □

Proof of Theorem 3.10. Assume first that $\mathcal{K}$ has DAP. Let $A \subseteq M_{\mathcal{K}}$ be finite. We have $A' = \langle A \rangle \subseteq \mathrm{ACL}(A)$, since A' is fixed pointwise by $\mathrm{Aut}(M_{\mathcal{K}}/A)$. Now let $b \in \mathrm{ACL}(A)$, with $|\,\mathrm{orb}_A(b)| = n$. Assume for contradiction that $b \notin A'$. Let $B = \langle A' \cup \{b\} \rangle$, and let $f\colon A' \hookrightarrow B$ be the inclusion. By repeatedly applying DAP, we can find a structure C containing $n+1$ isomorphic copies of B, $B_1, \dots, B_{n+1}$ with a common copy of A', which are disjoint over A'. Each B_i contains an element b_i corresponding to b, and disjointness implies that $b_i \neq b_j$ when $i \neq j$. By richness, we can embed C into $M_{\mathcal{K}}$ over A'.

Now for each $1 \le i \le n+1$, there is an isomorphism $f_i\colon B \to B_i$ fixing A pointwise and with $f(b) = b_i$. By homogeneity, these f_i lift to automorphisms $\sigma_i \in \mathrm{Aut}(M_{\mathcal{K}}/A)$. Thus $b_1, \dots, b_{n+1} \in \mathrm{orb}_A(b)$, contradicting $|\,\mathrm{orb}_A(b)| = n$.

Conversely, assume $\mathrm{ACL}(A) = \langle A \rangle$ for all $A \subseteq M_{\mathcal{K}}$ finite. Suppose we have $f_1\colon A \hookrightarrow B_1$ and $f_2\colon A \hookrightarrow B_2$ with $A, B_1, B_2 \in \mathcal{K}$. By AP and richness of $M_{\mathcal{K}}$, there exist embeddings $g_1\colon B_1 \hookrightarrow M_{\mathcal{K}}$ and $g_2\colon B_2 \hookrightarrow M_{\mathcal{K}}$ such that $g_1 \circ f_1 = g_2 \circ f_2$. Let $A' = g_1(f_1(A)) = g_2(f_2(A))$. The problem is that we may not have $g_1(B_1) \cap g_2(B_2) = A'$.

By assumption, letting $G = \mathrm{Aut}(M_{\mathcal{K}}/A')$, each element of $B_1 \setminus A'$ is not in $\mathrm{ACL}(A')$, so it has infinite orbit under the action of G. By Lemma 3.12, there exists $\sigma \in G$ such that $\sigma(B_1 \setminus A') \cap B_2 = \varnothing$, so $\sigma(B_1) \cap B_2 = A'$. Setting $g_1' = \sigma \circ g_1$ and $C = \langle g_1'(B_1) \cup g_2(B_2) \rangle$ shows that $\mathcal{K}$ has DAP. □

Lecture 4: Order properties

Model theorists love to measure how much order is present in a structure. The interest in this probably originates with Shelah's identification of the important class of stable theories and his characterization of stability by the absence of the order property. Fraïssé limits tend to be unstable (i.e., they usually have the order property, and even a stronger property called the independence property). But our goal in this section is to show that free amalgamation in a Fraïssé class $\mathcal{K}$ is incompatible with a certain strengthening of the order property in the theory of the Fraïssé limit $\mathrm{Th}(M_{\mathcal{K}})$.

Until further notice, T is a complete theory, and $\mathcal{U} \models T$ is a sufficiently saturated model. We write x and y for tuples of variables of arbitrary length. When we write $\varphi(x; y)$, we mean the free variables of the formula φ have been partitioned into two tuples, x and y.

Definition 4.1. T has the **order property** (OP) if there is a formula $\varphi(x; y)$ and sequences $(a_i)_{i\in\omega}$ from $\mathcal{U}^x$ and $(b_i)_{i\in\omega}$ from $\mathcal{U}^y$ such that $\mathcal{U} \models \varphi(a_i; b_j)$ if and only if $i < j$. Otherwise, T is NOP, or **stable**.

Definition 4.2. T has the **independence property** (IP) if there is a formula $\varphi(x;y)$ and sequences $(a_i)_{i\in\omega}$ from $\mathcal{U}^x$ and $(b_X)_{X\subseteq\omega}$ from $\mathcal{U}^y$ such that $\mathcal{U} \models \varphi(a_i;b_X)$ if and only if $i \in X$. Otherwise, T is NIP.

Note that if T has IP, witnessed by $\varphi(x;y)$, then the same formula witnesses OP. Indeed, letting $[j] = \{0,\ldots,j-1\}$, we have $\mathcal{U} \models \varphi(a_i;b_{[j]})$ if and only if $i \in [j]$ if and only if $i < j$.

Example 4.3. The complete theory $T^*_{\mathcal{K}_G}$ of the random graph has IP, and hence also OP, witnessed by xRy. Indeed, we can take $(a_i)_{i\in\omega}$ to be any sequence of distinct elements from $\mathcal{U}$. Then for each $X \subseteq \omega$, the partial type

$$p_X(y) = \{a_iRy \mid i \in X\} \cup \{\neg a_iRy \mid i \notin X\}$$

is consistent by the extension properties and compactness. We can take b_X to be any realization of $p_X(y)$ in $\mathcal{U}$.

Lemma 4.4. *T has* OP *if and only if there is a formula $\psi(z;z')$ with z and z' tuples of the same length, and a sequence $(c_i)_{i\in\omega}$ such that $\mathcal{U} \models \varphi(c_i;c_j)$ if and only if $i < j$.*

Proof. Such a $\psi(z;z')$ and $(c_i)_{i\in\omega}$ witnesses OP. Conversely, suppose $\varphi(x;y)$ and $(a_i)_{i\in\omega}$ and $(b_i)_{i\in\omega}$ witness OP. Let $z = xy$ and $z' = x'y'$, let $c_i = a_ib_i$ for all $i \in \omega$, and let $\psi(z,z') = \varphi(x;y')$. Then $\mathcal{U} \models \psi(c_i,c_j)$ if and only if $\mathcal{U} \models \varphi(a_i;b_j)$ if and only if $i < j$. ☐

If T has OP, then some formula $\psi(z;z')$ orders an infinite sequence $(c_i)_{i\in\omega}$ from $\mathcal{U}^z$. But the formula ψ need not behave anything like an order on all of $\mathcal{U}^z$. The following stronger definition imposes a global ordering relation.

Definition 4.5. T has the **strict order property** (SOP) if there is a formula $\varphi(x;y)$ and a sequence $(b_i)_{i\in\omega}$ from $\mathcal{U}^y$ such that $\varphi(\mathcal{U};b_i) \subsetneq \varphi(\mathcal{U};b_{i+1})$ for all i. Otherwise, T is NSOP.

Note that if $\varphi(x;y)$ and $(b_i)_{i\in\omega}$ witnesses SOP, then the formula

$$\forall x\,(\varphi(x;y) \to \varphi(x;y')) \land \exists x\,(\neg\varphi(x;y) \land \varphi(x;y'))$$

defines a strict partial order $<_\varphi$ (an irreflexive and transitive relation) on all of $\mathcal{U}^y$ which has an infinite chain: $b_i <_\varphi b_{i+1}$ for all $i \in \omega$. It follows that T has OP, witnessed by the formula $y <_\varphi y'$ and $(b_i)_{i\in\omega}$.

Conversely, if $y < y'$ is a definable strict partial order on $\mathcal{U}^y$ with an infinite chain $(b_i)_{i\in\omega}$, then $(\mathcal{U} < b_i) \subsetneq (\mathcal{U} < b_{i+1})$ for all $i \in \omega$, so T has SOP.

Example 4.6. DLO $= \mathrm{Th}(\mathbb{Q};\leq)$ has SOP, and hence also OP, witnessed by $x < y$ and $(b_i)_{i\in\omega}$ any strictly increasing sequence.

Fact 4.7 (Shelah)**.** T has OP if and only if T has IP or T has SOP.

It is an exercise to show that DLO is NIP and $T^*_{\mathcal{K}_G}$ is NSOP.

Definition 4.8. Let $n \geq 3$. T has the **n-strong order property** (SOP_n) if there is a formula $\varphi(x; x')$ (where x and x' have the same length) and a sequence $(a_i)_{i\in\omega}$ from $\mathcal{U}^x$ such that $\mathcal{U} \models \varphi(a_i; a_j)$ for all $i < j$, but there are no φ-cycles of length n, i.e., the following formula is inconsistent:

$$\bigwedge_{i<(n-1)} \varphi(x_i; x_{i+1}) \wedge \varphi(x_{n-1}; x_0).$$

Note that if $<$ is a definable strict partial order on $\mathcal{U}^x$ with infinite chains, then $x < x'$ has SOP_n. It follows that SOP implies SOP_n for all $n \geq 3$. We will show below that SOP_n implies OP for all $n \geq 3$. It is a (not-so-easy) exercise to show that SOP_{n+1} implies SOP_n for all $n \geq 3$.

Example 4.9. Let H be the Henson graph, the Fraïssé limit of the class $\mathcal{K}_\triangle$ of all finite triangle-free graphs. $\mathrm{Th}(H)$ has SOP_3.

Let $\varphi(x_1, x_2, x_3; y_1, y_2, y_3)$ be $x_1 R y_2 \wedge x_2 R y_3 \wedge x_3 R y_1$. Consider the countable graph G with vertices $\{a_i, b_i, c_i \mid i \in \omega\}$ and edges $a_i R b_j$ when $i < j$, $b_i R c_j$ when $i < j$, and $c_i R a_j$ when $i < j$ (and no other edges). I claim that G has no triangles. Indeed, since there are no edges $a_i R a_j$ or $b_i R b_j$ or $c_i R c_j$, any triangle must contain vertices a_i, b_j, c_k for some $i, j, k \in \omega$. But then $i < j$, $j < k$, and $k < i$, contradiction.

By universality of H, there is an embedding $G \hookrightarrow H$, and we identify G with its image in H. For all $i < j$, we have $H \models \varphi(a_i, b_i, c_i; a_j, b_j, c_j)$ by construction. But $\varphi(x_1, x_2, x_3; y_1, y_2, y_3) \wedge \varphi(y_1, y_2, y_3; z_1, z_2, z_3) \wedge \varphi(z_1, z_2, z_3; x_1, x_2, x_3)$ is inconsistent. Indeed, this conjunction implies $x_1 R y_2 \wedge y_2 R z_3 \wedge z_3 R x_1$, which is inconsistent, since H is triangle-free.

It will follow from Theorem 5.10 below that $\mathrm{Th}(H)$ is NSOP_4.

Example 4.10. Let $\mathcal{L} = \{D_q \mid q \in \mathbb{Q} \cap [0, 1]\}$ and $\mathcal{K}$ =finite metric spaces with all distances in $\mathbb{Q} \cap [0, 1]$. We view these metric spaces as $\mathcal{L}$-structures by setting $D_q(x, y)$ if and only if $d(x, y) = q$. Then $\mathcal{K}$ is a Fraïssé class, whose Fraïssé limit $M_\mathcal{K}$ is called the rational Urysohn sphere. As usual, $M_\mathcal{K}$ is characterized by one-point extension axioms: For every finite subspace $X = \{x_1, \ldots, x_n\}$ and every assignment of distances $d(x_i, y)$ in $\mathbb{Q} \cap [0, 1]$ which do not violate the triangle inequality, there exists a y with the specified distances to all the x_i.

The metric completion of $M_\mathcal{K}$ is the Urysohn sphere, the unique universal and homogeneous separable metric space with all distances ≤ 1.

It is a theorem of Conant and Terry that $\mathrm{Th}(M_\mathcal{K})$ is NSOP and has SOP_n for all $n \geq 3$. But it is necessary to use a different formula to witness SOP_n for each n.

Question 4.11. Is there an NSOP theory T with a single formula $\varphi(x; y)$ that witnesses SOP_n for all $n \geq 3$, i.e., such that there is a sequence $(a_i)_{i\in\omega}$ with $\mathcal{U} \models \varphi(a_i; a_j)$ for all $i < j$, but there are no φ-cycles of any finite length?

Lecture 5: Indiscernibles, free amalgamation

To prove things about these order properties, it is useful to make the witnesses indiscernible.

Definition 5.1. Let $(I, \leq)$ be a linear order, let $\mathcal{I} = (a_i)_{i \in I}$ be a sequence from $\mathcal{U}^x$, and let B be a set (when no set B is mentioned, we assume $B = \varnothing$). We say that $\mathcal{I}$ is a B**-indiscernible sequence** if for all $n \in \omega$, all $i_1 < \cdots < i_n$ and $j_1 < \ldots j_n$ in I, and all $\mathcal{L}(B)$-formulas $\varphi(x_1, \ldots, x_n)$ (where each x_i is a tuple of the same length as x), we have

$$\mathcal{U} \models \varphi(a_{i_1}, \ldots, a_{i_n}) \text{ iff } \models \varphi(a_{j_1}, \ldots, a_{j_n}).$$

Example 5.2. Let $T = \mathrm{DLO}$. By quantifier elimination, every strictly increasing sequence $(a_i)_{i \in \omega}$ of elements of $\mathcal{U}$ is an indiscernible sequence.

On the other hand, with $T = \mathrm{Th}(\mathbb{Z}; <)$ the strictly increasing sequence $0, 1, 2, \ldots$ is not indiscernible: letting $\varphi(x_1, x_2)$ be $\exists y\, (x_1 < y \wedge y < x_2)$, we have $\models \neg\varphi(0, 1)$, but $\models \varphi(0, 2)$.

Let $T = T^*_{\mathcal{K}_G}$. By quantifier elimination, every infinite clique or independent set is an indiscernible sequence.

Definition 5.3. Let $(I, \leq)$ be an infinite linear order, let $\mathcal{I} = (a_i)_{i \in I}$ be an I-indexed sequence from $\mathcal{U}^x$ (not necessarily indiscernible), and let B be a set. The **Ehrenfeucht–Mostowski type** of $\mathcal{I}$ over B is a set of $\mathcal{L}(B)$-formulas in contexts $x_1, \ldots, x_n$, where $n \in \omega$ and each x_i is a tuple of the same length as x:

$$\mathrm{EM}(\mathcal{I}/B) = \{\varphi(x_1, \ldots, x_n) \in \mathcal{L}(B) \mid \text{for all } i_1 < \cdots < i_n \in I, \models \varphi(a_{i_1}, \ldots, a_{i_n})\}.$$

A sequence $\mathcal{J} = (a'_i)_{j \in J}$ from $\mathcal{U}^x$, indexed by a linear order $(J, \leq)$, **satisfies** $\mathrm{EM}(\mathcal{I}/B)$ if for all $\varphi(x_1, \ldots, x_n) \in \mathrm{EM}(\mathcal{I}/B)$, we have $\models \varphi(a'_{j_1}, \ldots, a'_{j_n})$ for all $j_1 < \cdots < j_n \in J$. We also say that $\mathcal{J}$ is **locally based** on $\mathcal{I}$ over B.

The sequence $\mathcal{I}$ is indiscernible over B if and only if $\mathrm{EM}(\mathcal{I}/B)$ is complete in the sense that for any L_B-formula $\varphi(x_1, \ldots, x_n)$, either $\varphi \in \mathrm{EM}(\mathcal{I}/B)$ or $\neg\varphi \in \mathrm{EM}(\mathcal{I}/B)$.

Using Ramsey's Theorem, it is always possible to take a sequence in a model of T and find an indiscernible sequence locally based on it.

Lemma 5.4 ("Standard Lemma")**.** *Let $\mathcal{I} = (a_i)_{i \in \omega}$ be a sequence from $\mathcal{U}^x$, and let B be a set. Then there is a B-indiscernible sequence $\mathcal{J} = (c_j)_{j \in \omega}$ satisfying* $\mathrm{EM}(\mathcal{I}/B)$.

Proof. Consider the partial type q (in context $(y_j)_{j \in \omega}$, where each y_j is a tuple of length $|x|$) consisting of formulas:

(a) $\varphi(y_{j_1}, \ldots, y_{j_n})$, where $j_1 < \cdots < j_n$ in ω and $\varphi(x_1, \ldots, x_n) \in \mathrm{EM}(\mathcal{I}/B)$.

(b) $\varphi(y_{j_1}, \ldots, y_{j_n}) \leftrightarrow \varphi(y_{j'_1}, \ldots, y_{j'_n})$, where $j_1 < \cdots < j_n$ and $j'_1 < \cdots < j'_n$ in ω and $\varphi(x_1, \ldots, x_n)$ is an L_B-formula.

It suffices to show that q is consistent, since the formulas of type (a) ensure that $\mathcal{J}$ satisfies $\mathrm{EM}(\mathcal{I}/B)$ and the formulas of type (b) ensure that $\mathcal{J}$ is indiscernible over B.

For any finite set of $\mathcal{L}(B)$-formulas Δ, let q_Δ be the same partial type, but with the formulas of type (b) restricted to those $\mathcal{L}(B)$-formulas appearing in Δ. A finite subset of q is contained in q_Δ for some finite set Δ, so by compactness it suffices to show that q_Δ is consistent. Further, by adding dummy variables, we may assume that each formula in Δ has the same context $x_1, \ldots, x_n$. Indeed, if $\psi(x_1, \ldots, x_m)$ is an $\mathcal{L}(B)$-formula with $m < n$, let $\psi'(x_1, \ldots, x_n)$ be the same formula with dummy variables $x_{m+1}, \ldots, x_n$ added. Now for any $j_1 < \cdots < j_m$ and $j'_1 < \cdots < j'_m$ in ω, we can extend both sequences to $j_1 < \cdots < j_n$ and $j'_1 < \cdots < j'_n$ in ω, and (b) for ψ' implies

$$\psi(y_{j_1}, \ldots, y_{j_m}) \leftrightarrow \psi'(y_{j_1}, \ldots, y_{j_n}) \leftrightarrow \psi'(y_{j'_1}, \ldots, y_{j'_n}) \leftrightarrow \psi(y_{j'_1}, \ldots, y_{j'_m}).$$

Let $[\omega]^n$ be the set of all strictly increasing n-tuples from ω. Define a coloring $c\colon [\omega]^n \to \mathcal{P}(\Delta)$ by

$$c(\{i_1, \ldots, i_n\}) = \{\varphi(x_1, \ldots, x_n) \in \Delta \mid \mathcal{U} \models \varphi(a_{i_1}, \ldots, a_{i_n})\}.$$

By Ramsey's Theorem, there is an infinite subset $H \subseteq \omega$ which is homogeneous for c. Enumerate H in increasing order as $(a_{i_j})_{j\in\omega}$. This sequence satisfies q_Δ, which completes the proof. □

To illustrate the use of making witnesses to order properties indiscernible, consider the following.

Lemma 5.5. *If T has SOP_n for some $n \geq 3$, then T has* OP.

Proof. Suppose $\varphi(x; x')$ and $(a_i)_{i\in\omega}$ witness SOP_n. We have $\mathcal{U} \models \varphi(a_i; a_j)$ for all $i < j$, and we would be done if we could show that for all $i \geq j$, $\mathcal{U} \not\models \varphi(a_i; a_j)$. This should be ruled out by the inconsistency of cycles, but we could have $\mathcal{U} \models \varphi(a_i; a_j)$ with $i \geq j$ but i and j "too close" to form an n-cycle.

By Lemma 5.4, we can find $(a'_i)_{i\in\omega}$ indiscernible and locally based on $(a_i)_{i\in\omega}$. Since $\mathcal{U} \models \varphi(a_i; a_j)$ for all $i < j$, $\varphi(x; x') \in \mathrm{EM}((a_i)_{i\in\omega}/\varnothing)$, so $\mathcal{U} \models \varphi(a'_i; a'_j)$ for all $i < j$. Suppose for contradiction that we have $\mathcal{U} \models \varphi(a'_i; a'_j)$ for some $i > j$. By indiscernibility, $\mathcal{U} \models \varphi(a'_{n-1}; a'_0)$. But then $(a'_i)_{i<n}$ forms a φ-cycle of length n, contradiction.

By indiscernibility, either $\mathcal{U} \models \neg\varphi(a'_i; a'_i)$ for all i, or $\mathcal{U} \models \varphi(a'_i; a'_i)$ for all i. In the first case, we have a witness to OP. In the second case, $\varphi(x; x')$ and $(a'_{i+1})_{i\in\omega}$ and $(a'_i)_{i\in\omega}$ witnesses OP, since $\mathcal{U} \models \varphi(a'_{i+1}; a'_j)$ if and only if $i+1 \leq j$ if and only if $i < j$. □

We need one more useful lemma about indiscernibles.

Lemma 5.6. *Let $(a^i)_{i\in I}$ be a B-indiscernible sequence, indexed by an infinite linear order I, with each a^i an n-tuple $(a^i_0, \ldots, a^i_{n-1})$. Then there exists $X \subseteq [n] = \{0, \ldots, n-1\}$ such that for all $k \in X$, $a^i_k = a^j_k$ for all $i \neq j$ in I, and* *for all $k \notin X$ and $k' \in [n]$, $a^i_k \neq a^j_{k'}$ for all $i \neq j$ in I. In other words, each a^i can be partitioned into subtuples b^i (the coordinates not in X) and c^i (the coordinates in X) such that $c^i = c^j$ for all $i \neq j$ and b^i and b^j are disjoint for all $i \neq j$.*

Proof. Let X be the set of all $k \in [n]$ such that there exist $i < j$ in I with $a^i_k = a^j_k$. By indiscernibility, for all $k \in X$, for all $i \neq j$ in I, $a^i_k = a^j_k$.

Now let $k \notin X$ and $k' \in [n]$. Let $i \neq j$ in I, and suppose for contradiction that $a^i_k = a^j_{k'}$. Assume $i < j$ (the other case is similar), and since I is infinite (shifting i or j if necessary), we may assume without loss of generality that there is some i' in I with $i < i' < j$. Since $a^i_k = a^j_{k'}$, by indiscernibility $a^{i'}_k = a^j_{k'}$. But then $a^i_k = a^{i'}_k$, so $k \in X$, contradiction. □

For the remainder of the section, we assume the language $\mathcal{L}$ is relational.

We now turn to our Fraïssé-theoretic sufficient condition for NSOP_4. The proof is implicit in the verifications that several examples are NSOP_4 (by Shelah and others), but it was given in its general form for the first time (a bit more generally than we present here) by Rehana Patel in unpublished work. More recently, Gabe Conant abstracted the proof away from the setting of Fraïssé limits, to theories equipped with a certain kind of stationary independence relation called a "free amalgamation relation". Scott Mutchnik has further developed these ideas in recent work.

For sets A, B, C, and D, we write $D = B \sqcup_A C$ to indicate that D is the disjoint union of B and C over A: $B \cup C = D$ and $B \cap C = A$.

Definition 5.7. Let C be a structure with substructures A, B_1, and B_2 such that $A \subseteq B_1$ and $A \subseteq B_2$. We say that C is the **free amalgam** of B_1 and B_2 over A, denoted $C = B_1 \oplus_A B_2$, if $C = B_1 \sqcup_A B_2$, and for every relation symbol $R \in \mathcal{L}$, $R(C) = R(B_1) \sqcup_{R(A)} R(B_2)$. More explicitly, if $R(\overline{c})$ holds in C, then $\overline{c} \in B_1^n$, or $\overline{c} \in B_2^n$.

Definition 5.8. We say a Fraïssé class $\mathcal{K}$ has **free amalgamation** if for all $f_1\colon A \hookrightarrow B_1$ and $f_2\colon A \hookrightarrow B_2$ in $\mathcal{K}$, there exists $C \in \mathcal{K}$ and $g_1\colon B_1 \hookrightarrow C$ and $g_2\colon B_2 \hookrightarrow C$ such that $g_1 \circ f_1 = g_2 \circ f_2$, and furthermore, letting $A' = g_1(f_1(A)) = g_2(f_2(A)) \subseteq C$ we have $C = g_1(B_1) \oplus_{A'} g_2(B_2)$.

Note that free amalgamation is stronger than disjoint amalgamation. There is not an obvious way to extend the definition of free amalgamation to languages with function symbols. One possibility would be to require C to be the *free $\mathcal{L}$-structure* generated by B_1 and B_2 over A. But we will stick to the relational context for simplicity.

Example 5.9. The class $\mathcal{K}_{\mathrm{G}}$ of finite graphs and the class $\mathcal{K}_{\triangle}$ of finite triangle-free graphs both have free amalgamation. The class $\mathcal{K}_{\mathrm{LO}}$ does not have free amalgamation, since for any $b_1 \in B_1 \setminus A$ and $b_2 \in B_2 \setminus A$, we must have either $b_1 \leq b_2$ or $b_2 \leq b_1$ in the amalgam C.

Now we're ready to prove our main Theorem. Example 4.9 shows the result is sharp: we cannot obtain NSOP_3 in general.

Theorem 5.10. *Suppose $\mathcal{K}$ is a finitary Fraïssé class with free amalgamation. Let $T = \mathrm{Th}(M_{\mathcal{K}})$, the theory of the Fraïssé limit. Then T is NSOP_4, and hence NSOP.*

Proof. Suppose for contradiction that $\varphi(x; x')$ and $(a_i)_{i\in\omega}$ witnesses SOP_4. Exactly as in Lemma 5.5, we may assume that $(a_i)_{i\in\omega}$ is indiscernible. Since $(a_i)_{i\in\omega}$ is countable, by Löwenheim–Skolem, it lives in a countable model of T, and by $\aleph_0$-categoricity (Theorem 3.7), in $M_{\mathcal{K}}$.

By Lemma 5.6, we can divide each a_i into subtuples b_i and c_i, where for all $i < j$ in ω, $c_i = c_j$ and b_i is disjoint from b_j.

Let $C \in \mathcal{K}$ be the substructure enumerated by c_0, and for each $i \in \omega$, let $A_i \in \mathcal{K}$ be the substructure enumerated by a_i. All A_i have C as a common substructure, and by indiscernibility, all A_i are isomorphic by isomorphisms mapping $a_i \mapsto a_j$ (and hence fixing C pointwise).

Let $B_0 = A_0 \cup A_1$, and let $B_2 = A_1 \cup A_2$. Let $D = B_0 \oplus_{A_1} B_2$. By free amalgamation, $D \in \mathcal{K}$, and by richness of $M_{\mathcal{K}}$, there is an embedding $f\colon D \hookrightarrow M_{\mathcal{K}}$ over B_2. Let $D' = f(D)$, $B_0' = f(B_0)$ and $A_0' = f(A_0)$. So $D' = B_0' \oplus_{A_1} B_2'$. We also write a_0' for the tuple which is the image of a_0 under this embedding.

Let $E = A_0' \cup A_2$. Since $A_0' \subseteq B_0'$, $A_2 \subseteq B_2$, and $A_0' \cap A_2 = C$, we have $E = A_0' \oplus_C A_2$. Now E has an automorphism swapping A_0' and A_2 and fixing C pointwise. By homogeneity, there exists an automorphism $\sigma \in \mathrm{Aut}(M_{\mathcal{K}})$ with $\sigma(a_0') = a_2$ and $\sigma(a_2) = a_0'$. Let $a_1' = \sigma(a_1)$.

By homogeneity, since $B_0 \cong B_0'$, $\mathrm{tp}(a_0'a_1) = \mathrm{tp}(a_0a_1)$, and in particular $M_{\mathcal{K}} \models \varphi(a_0', a_1)$. By hypothesis, $M_{\mathcal{K}} \models \varphi(a_1, a_2)$. Then also $M_{\mathcal{K}} \models \varphi(\sigma(a_0'), \sigma(a_1))$ and $M_{\mathcal{K}} \models \varphi(\sigma(a_1), \sigma(a_2))$, so $M_{\mathcal{K}} \models \varphi(a_2, a_1')$ and $M_{\mathcal{K}} \models \varphi(a_1', a_0')$. Thus a_0', a_1, a_2, a_1' form a φ-cycle of length 4, contradicting SOP_4. □

Lecture 6: Zero-one laws and pseudofiniteness

We now turn to another notion of limit of finite structures, the *logical limit*, with an interest in comparing it with the Fraïssé limit. Logic limits in the sense of first-order zero-one laws were introduced independently by Glebskii, Kogan, Liogon'kii, and Talanov (in the USSR in 1969) and Fagin (in the US in 1976), and they are now a core area of finite model theory. For a survey, see Compton "0-1 Laws in Logic and Combinatorics". It is possible to study zero-one laws for other logics, but we stick to the first-order case here.

Let $\mathcal{L}$ be a relational language. Given a class $\mathcal{K}$ of finite structures, let $\mathcal{K}(n)$ be the set of structures in $\mathcal{K}$ with domain $[n]$. We assume $\mathcal{K}$ is finitary, has HP, and contains arbitrarily large finite structures, so $\mathcal{K}(n)$ is finite and non-empty for all n. Note that $\mathcal{K}(n)$ may contain multiple isomorphic copies of a single structure, with different labelings of the domain by $[n]$. In fact, the number of times a structure A appears (up to isomorphism) in $\mathcal{K}(n)$ is $n!/|\mathrm{Aut}(A)|$. It is also possible to consider zero-one laws for unlabeled structures, but the counting is usually easier for labeled structures.

For each n, let μ_n be a probability measure on $\mathcal{K}(n)$. For a property P, we write $[P] = \{A \in \mathcal{K}(n) \mid A \text{ satisfies } P\}$ (usually P is a formula $\varphi(\overline{a})$, where $\overline{a}$ is a tuple from $[n]$). We are often interested in the **uniform measure**:

$$\mu_n([P]) = \frac{|[P]|}{|\mathcal{K}(n)|}.$$

Given a sentence φ, and a sequence $\mu = (\mu_n)_{n\in\omega}$ of probability measures on the $\mathcal{K}(n)$, the **limit probability** of φ is

$$\mu_\infty(\varphi) = \lim_{n\to\infty} \mu_n([\varphi]).$$

We denote by $T^\mu_\mathcal{K}$ the **almost sure theory**: the set of sentences with limit probability 1. We say $\mathcal{K}$ has a **first-order zero-one law** with respect to μ if every sentence φ has limit probability 0 or 1. Since $\mu_\infty(\neg\varphi) = 1 - \mu_\infty(\varphi)$, $T^\mu_\mathcal{K}$ has a first-order zero-one law if and only if $T^\mu_\mathcal{K}$ is complete.

Lemma 6.1. *Let Σ be a countable set of sentences. Suppose that $\mu_\infty(\varphi) = 1$ for all $\varphi \in \Sigma$. Then if $\Sigma \models \psi$, we have $\mu_\infty(\psi) = 1$.*

Proof. By compactness, there are finitely many sentences $\varphi_0, \ldots, \varphi_{m-1} \in \Sigma$ such that $\bigwedge_{i<m} \varphi_i \models \psi$. For any $\varepsilon > 0$, since $\mu_\infty(\varphi) = 1$, there exists N_ε such that for all $n \geq N_\varepsilon$, $\mu_n(\varphi_i) > 1 - \frac{\varepsilon}{m}$ for all $i < m$. Then $\mu_n(\bigwedge_{i<m} \varphi_i) > (1-\varepsilon)$, so $\mu_n(\psi) > 1 - \varepsilon$. Since ε was arbitrary, $\mu_\infty(\psi) = 1$. □

Theorem 6.2. *The class $\mathcal{K}_\mathrm{G}$ of finite graphs has a first-order zero-one law with respect to the uniform measures μ. Moreover, $T^\mu_{\mathcal{K}_\mathrm{G}} = T^*_{\mathcal{K}_\mathrm{G}}$, the theory of the random graph.*

Proof. It may be helpful to observe that the uniform measure μ_n on $\mathcal{K}_\mathrm{G}(n)$ is equivalent to the following random construction (called the Erdös–Renyi random graph model $G(n, \frac{1}{2})$): for all $i < j < n$, flip a fair coin. If it is heads, put an edge between vertices i and j. Indeed, exactly half the graphs in $\mathcal{K}_\mathrm{G}(n)$ have an edge between i and j, and these edge probabilities are independent for distinct edges.

Recall that $T^*_{\mathcal{K}_\mathrm{G}}$ is axiomatized by the universal theory of graphs, together with one-point extension axioms for $0 \leq k \leq n$:

$$\varphi_{n,k} \colon \forall x_1, \ldots, x_n \left(\bigwedge_{i<j<n} x_i \neq x_j \to \exists y \left(\bigwedge_{i<k} yRx_i \wedge \bigwedge_{k\leq i<n} \neg yRx_i \right)\right).$$

By Lemma 6.1, if we show that each of these axioms has limiting probability 1, then $T^*_{\mathcal{K}_\mathrm{G}} \subseteq T^\mu_{\mathcal{K}_\mathrm{G}}$. Since $T^*_{\mathcal{K}_\mathrm{G}}$ is complete (Theorem 3.7), and $T^\mu_{\mathcal{K}_\mathrm{G}}$ is consistent, it follows that they are equal.

Since every structure in $\mathcal{K}_G(n)$ is a graph, the graph axioms have limiting probability 1. So we consider an extension axiom $\varphi_{n,k}$. For $N \in \omega$, we compute the probability that $\varphi_{n,k}$ is *not* satisfied by a graph on $[N]$.

Let $\theta_{n,k}(\overline{x}, y)$ be the formula $\left(\bigwedge_{i<k} yRx_i \wedge \bigwedge_{k\leq i<n} \neg yRx_i\right)$. Given a tuple $\overline{a}$ of distinct elements in $[N]$ and b not in $\overline{a}$, $\mu_N([\theta_{n,k}(\overline{a}, b)]) = 2^{-n}$, and for $b \neq b'$ in $[N]$ (and not in $\overline{a}$), the events $[\theta_{n,k}(\overline{a}, b)]$ and $[\theta_{n,k}(\overline{a}, b')]$ are independent. Thus, since there are $N - n$ choices for the witness b, we have

$$\mu_N([\forall y \, \neg\theta(\overline{a}, y)]) = (1 - 2^{-n})^{N-n}.$$

Now there are N^n choices for the tuple $\overline{a}$ from $[N]$ (and the redundant tuples contribute nothing to the probability), so we have an upper bound:

$$\mu_N([\exists \overline{x} \, (\bigwedge_{i<j<n} x_i \neq x_j \wedge \forall y \, \neg\theta(\overline{a}, y))]) \leq N^n(1 - 2^{-n})^{N-n}.$$

Since n is constant, the exponential decay term $(1 - 2^{-n})^{N-n}$ dominates the polynomial growth term N^n in the limit $N \to \infty$, so $\mu_\infty(\neg\varphi_{n,k}) = 0$, and hence $\mu_\infty(\varphi_{n,k}) = 1$. □

As a consequence of the theorem, we determine which (first-order definable!) properties hold of "almost all" finite graphs by answering the same question for the random graph. For example, almost all finite graphs are connected, in fact of diameter 2, since the random graph satisfies $\forall x \forall y \exists z \, (xRz \wedge zRy)$. Of course, first-order logic over finite graphs is not expressive enough to define many properties of interest, which is why there is significant interest in extending zero-one laws to stronger logics.

Essentially the same proof as for Theorem 6.2 shows that if $\mathcal{K}_\mathcal{L}$ is the class of all finite $\mathcal{L}$-structures in a finite relational language $\mathcal{L}$, then $\mathcal{K}_\mathcal{L}$ has a first-order zero-one law with respect to the uniform measures, and $T^\mu_{\mathcal{K}_\mathcal{L}}$ is the theory of the Fraïssé limit of $\mathcal{K}_\mathcal{L}$ (this is the case originally considered by Fagin). It is crucial here that $\mathcal{L}$ is relational. For example, if $\mathcal{L}$ contains a single unary function symbol f, then

$$\mu_\infty(\forall x \, (f(x) \neq x)) = \frac{1}{e},$$

so $\mathcal{K}_\mathcal{L}$ does not have a first-order zero-one law.

An immediate consequence of Theorem 6.2 is that the theory $T^*_{\mathcal{K}_G}$ is pseudofinite.

Definition 6.3. A theory T is **pseudofinite** if for every sentence φ such that $T \models \varphi$, φ has a finite model.

Example 6.4. If a theory T is *not* pseudofinite, then it entails a sentence which has only infinite models. Here are two standard ways that this can happen:

1. Unbounded orders: The conjunction of the strict partial order axioms with $\forall x \exists y \, x < y$ is a sentence that has only infinite models. For example, $\mathrm{DLO} = \mathrm{Th}(\mathbb{Q}, \leq)$ is not pseudofinite.

2. Definable functions which are injective but not surjective (or vice versa). For example, $\mathrm{Th}(\mathbb{Z},+)$ is not pseudofinite, because the function $x \mapsto x+x$ is injective but not surjective. And $\mathrm{Th}(\overline{\mathbb{F}_p})$ (for $p \neq 2$) is not pseudofinite, because the function $x \mapsto x \cdot x$ is surjective but not injective.

Returning to graphs, if $T^*_{\mathcal{K}_\mathrm{G}} \models \varphi$, we have seen not just that φ has a finite model, but that *almost all* sufficiently large graphs satisfy φ. Of course, this is a probabilistic argument, not a particularly constructive one. Can we actually find a finite graph satisfying φ in a better way than enumerating all sufficiently large graphs and checking? It turns out the answer is yes.

Example 6.5. Given a prime number p congruent to 1 mod 4, we define a graph G_p with domain the finite field $\mathbb{F}_p$ by setting aRb if and only if $(b-a)$ is a non-zero square in $\mathbb{F}_p$. Since $p \equiv 1 \mod 4$, -1 is a square in $\mathbb{F}_p$, so this relation is symmetric. Any sentence in the theory of the random graph is true in G_p for all sufficiently large p.

The graphs G_p are called **Paley graphs**, and they are discrete models for the Erdös–Renyi random graph model $G(n, \frac{1}{2})$ in the sense that as $p \to \infty$, the densities of all edge configurations converge to the probabilities specified by $G(n, \frac{1}{2})$. For this reason, they are called "quasi-random" graphs. The first analysis of the Paley graphs used very non-trivial number theory. Chung, Graham, and Wilson showed that finitely many instances of limiting densities suffice to imply the rest, and these finitely many instances can be checked in an elementary way for the Paley graphs.

Theorem 6.2 may lead one to think that the equality $T^\mu_{\mathcal{K}_\mathrm{G}} = T^*_{\mathcal{K}_\mathrm{G}}$ and pseudofiniteness of $T^*_{\mathcal{K}_\mathrm{G}}$ are "typical" of Fraïssé limits, at least in finite relational languages. This is far from the case.

Example 6.6. Let $\mathcal{K}_\mathrm{LO}$ be the class of finite linear orders. $T^*_{\mathcal{K}_\mathrm{LO}}$ is DLO $= \mathrm{Th}(\mathbb{Q}, \leq)$. But this theory is not pseudofinite, as explained in Example 6.4 above. In fact, $\mathcal{K}_\mathrm{LO}$ has a zero-one law for the uniform measures, and $T^\mu_{\mathcal{K}_\mathrm{LO}}$ is the theory of infinite discrete linear orders with endpoints. Here "discrete" means that if x is not the greatest element, then it has an immediate successor, and if x is not the least element, then it has an immediate predecessor.

Example 6.7. Let $\mathcal{K}_\triangle$ be the class of finite triangle-free graphs. It is a theorem of Erdős, Kleitman, and Rothschild that almost all triangle-free graphs are bipartite. That is, with respect to the uniform measures μ, $T^\mu_{\mathcal{K}_\triangle}$ contains sentences asserting that there are no cycles of any odd length. This is in contrast to $T^*_{\mathcal{K}_\triangle} = \mathrm{Th}(H)$: since $\mathrm{Age}(H) = \mathcal{K}_\triangle$, H contains cycles of all odd lengths ≥ 5.

Kolaitis, Prömel, and Rothschild built on the work of Erdős–Kleitman–Rothschild and showed that $\mathcal{K}_\triangle$ has a zero-one law with respect to the uniform measures. $T^\mu_{\mathcal{K}_\triangle}$ is the theory of the generic bipartite graph: If we expand the language of graphs by two predicates P and Q for the sides of a bipartition, the class of finite bipartite graphs becomes a Fraïssé class. $T^\mu_{\mathcal{K}_\triangle}$ is the reduct to the graph language of the theory of the Fraïssé limit.

It is possible that $T^*_{\mathcal{K}_\triangle} = \mathrm{Th}(H)$ is pseudofinite, despite this not being witnessed by a zero-one law with respect to the uniform measures. In fact, this is a longstanding open problem.

Greg Cherlin has done a lot of hard combinatorial work on this problem. Michael Albert constructed an infinite family of finite triangle-free graphs satisfying the one-point extension axioms $\varphi_{A,B}$ with $|A| = 3$. But it is open whether there is any finite triangle-free graph satisfying the $\varphi_{A,B}$ with $|A| = 4$.

It seems likely to me that if $\mathrm{Th}(H)$ is pseudofinite, the finite models are sporadic, in the sense that they only occur in certain sizes or much have a very regular structure, in contrast to the wealth of finite models given to us by a zero-one law for the uniform measures.

The discussion above suggests the following very broad question:

Question 6.8. Which $\aleph_0$-categorical theories are pseudofinite?

Since every $\aleph_0$-categorical theory can be axiomatized by extension axioms in an expanded language (by Morleyization and Proposition 3.8), this question comes down to the question of when one-point extension axioms are satisfiable in finite structures. On the positive side, Cherlin, Harrington, and Lachlan proved that every $\aleph_0$-categorical and $\aleph_0$-stable theory is pseudofinite ($\aleph_0$-stability is a strengthening of stability, so such theories do not have the order property). On the negative size, we have the following:

Proposition 6.9. *No $\aleph_0$-categorical pseudofinite theory has the strict order property.*

Proof. If T has the strict order property, then it defines a strict partial order $<$ with infinite chains. Let $\{a_i \mid i \in \omega\}$ be a chain in the unique countable model M of T, so $M \models a_i < a_j$ if and only if $i < j$. In a countably categorical theory, automorphism-invariant properties are definable, so there is a formula $\varphi(x)$, with $\varphi(x) \in \mathrm{tp}(a_i)$ for all i, such that $M \models \varphi(b)$ if and only if there is an infinite $<$-chain above b.

Now $M \models \exists x\, \varphi(x) \wedge \forall x\, (\varphi(x) \to \exists y\, (x < y \wedge \varphi(y)))$. But in any partial order, this sentence implies the existence of an infinite increasing chain of elements satisfying $\varphi(x)$, so its conjunction with the partial order axioms for $<$ has no finite model. $\square$

Lecture 7: Higher amalgamation properties

In this section $\mathcal{L}$ is a relational language and $\mathcal{K}$ is a finitary Fraïssé class.

Before giving any formal definitions, let me start with some motivating discussion. The amalgamation property AP is about amalgamating two structures over a common substructure:

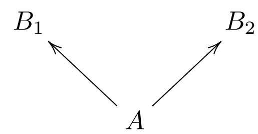

What if we want to amalgamate three structures? Since $\mathcal{K}$ has AP, it is easy to amalgamate a diagram like this:

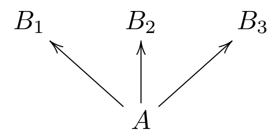

by amalgamating B_1 and B_2 and then amalgamating the result with B_3.

It's more interesting to specify amalgams of each of the pairs $\{B_1, B_2\}$, $\{B_1, B_3\}$, and $\{B_2, B_3\}$ and ask whether the resulting triple of structures can be amalgamated coherently:

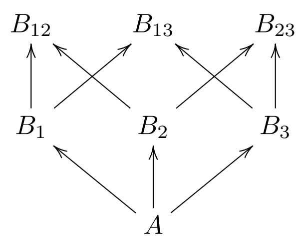

The answer is clearly no in general: if $b_i \in B_i \setminus A$ for all $1 \leq i \leq 3$, we might have $b_1 = b_2$ in B_{12} and $b_1 = b_3$ in B_{13}, but $b_2 \neq b_3$ in B_{23}. Then B_{12}, B_{13} and B_{23} cannot be embedded into a common structure C in such a way that all the squares in the diagram commute. Problems like this can arise whenever the images of B_i and B_j are not disjoint over A inside B_{ij}. So we will impose a disjointness requirement.

More generally, to amalgamate n structures, we will consider diagrams indexed by $\mathcal{P}^-([n])$, the set of all proper subsets of $[n]$.

Finally, it will be convenient to think about amalgamating types (relative to the complete theory $\mathrm{Th}(M_{\mathcal{K}})$), rather than structures. This is equivalent, since every complete type $p(x)$ relative to $\mathrm{Th}(M_{\mathcal{K}})$ is isolated by a formula $\theta_A(x)$ describing the isomorphism type of the induced substructure on x.

Let T be a theory. We say that a type $p(\overline{x})$ (with no parameters) in the variables $\{x_i \mid i \in I\}$ is **non-redundant** if it contains the formulas $\{x_i \neq x_j \mid i \neq j \in I\}$.

A family $\mathcal{F} \subseteq \mathcal{P}([n])$ of subsets of $[n]$ is **downwards closed** if $S' \in \mathcal{F}$ whenever $S' \subseteq S$ and $S \in \mathcal{F}$. Given a downwards closed family of subsets $\mathcal{F} \subseteq \mathcal{P}([n])$, and pairwise disjoint tuples of variables $\overline{x}_S$ for each $S \in \mathcal{F}$, a

coherent $\mathcal{F}$-family of types is a set $\{p_S \mid S \in \mathcal{F}\}$ such that each p_S is a non-redundant type in the variables $(\overline{x}_{S'})_{S' \subseteq S}$, and $p_{S'} \subseteq p_S$ when $S' \subseteq S$.

We denote by $\mathcal{P}^-([n])$ the downwards closed family of all proper subsets of $[n]$. For $n \geq 2$, a **disjoint n-amalgamation problem** is a coherent $\mathcal{P}^-([n])$-family of types. A **solution** to a disjoint n-amalgamation problem is an extension of the coherent $\mathcal{P}^-([n])$-family of types to a coherent $\mathcal{P}([n])$-family of types; that is, a non-redundant type $p_{[n]}$ such that $p_S \subseteq p_{[n]}$ for all S. We say T has the **disjoint n-amalgamation property (n-DAP)** if every n-amalgamation problem has a solution.

A **basic** disjoint n-amalgamation problem is a disjoint n-amalgamation problem such that each tuple of variables $\overline{x}_i$ is a singleton x_i, and all other tuples $\overline{x}_S$ with $|S| \neq 1$ is empty. So each p_S is a type in the variables $(x_i)_{i \in S}$. A **solution** to the basic disjoint n-amalgamation problem is a non-redundant type $p_{[n]}$ in the variables $(x_i)_{i \in [n]}$ such that $p_S \subseteq p_{[n]}$ for all S. Similarly, we say T has **basic n-DAP** if every basic disjoint n-amalgamation property has a solution.

Remark 7.1. A Fraïssé class $\mathcal{K}$ has DAP if and only if $T^*_{\mathcal{K}} = \mathrm{Th}(M_{\mathcal{K}})$ has 2-DAP. Given $A, B_0, B_1 \in \mathcal{K}$ and embeddings $f_0 \colon A \hookrightarrow B_0$ and $f_1 \colon A \hookrightarrow B_1$, let $\overline{x}_{\varnothing}$ be a tuple of variables enumerating A, and let $\overline{x}_{\{0\}}$ and $\overline{x}_{\{1\}}$ be tuples of variables enumerating $B_1 \setminus f_1(A)$ and $B_2 \setminus f_2(A)$. Set $p_{\varnothing} = \{\theta_A(\overline{x}_{\varnothing})\}$, and $p_{\{i\}} = \{\theta_{B_i}(\overline{x}_{\varnothing}, \overline{x}_{\{i\}})\}$. By QE, these determine complete types relative to $T^*_{\mathcal{K}}$. A solution to this disjoint 2-amalgamation problem is a non-redundant type $p_{\{0,1\}}(\overline{x}_{\varnothing}, \overline{x}_{\{0\}}, \overline{x}_{\{1\}}, \overline{x}_{\{0,1\}})$, which describes the isomorphism type of a structure C into which B_0 and B_1 embed disjointly over the image of A.

In general, we say that a finitary Fraïssé class $\mathcal{K}$ has n-DAP if the theory of its Fraïssé limit does.

Example 7.2. The class $\mathcal{K}_{\mathrm{G}}$ of finite graphs has n-DAP for all n. We can see this by generalizing free amalgamation: take $p_{[n]}$ to be the type saying that there are no edges other than those specified by the p_S for $S \in \mathcal{P}^-([n])$.

Example 7.3. Here are three examples of failures of basic 3-DAP in Fraïssé classes with 2-DAP:

1. Let $\mathcal{K}$ be the class of finite equivalence relations. The non-redundant 2-types specified by $\{x_0 E x_1\}$, $\{x_0 E x_2\}$, and $\{\neg x_1 E x_2\}$ cannot be amalgamated.

2. Let $\mathcal{K}_{\mathrm{LO}}$ be the class of finite linear orders. The non-redundant 2-types specified by $\{x_0 < x_1\}$, $\{x_1 < x_2\}$, and $\{x_2 < x_1\}$ cannot be amalgamated.

3. Let $\mathcal{K}_{\triangle}$ be the class of finite $\triangle$-free graphs. The non-redundant 2-types specified by $\{x_0 R x_1\}$, $\{x_0 R x_2\}$, and $\{x_1 R x_2\}$ cannot be amalgamated.

Example 7.4. Generalizing Example 7.3(3), let $\mathcal{K}^k_n$ be the class of n-free k-hypergraphs: the language consists of a single k-ary relation R, and the structures in $\mathcal{K}^k_n$ are k-hypergraphs with no substructure isomorphic to the complete hypergraph on n vertices. Note that $\mathcal{K}^2_3 = \mathcal{K}_{\triangle}$.

For $n > k$, $\mathcal{K}_n^k$ satisfies basic disjoint m-amalgamation for $m < n$, but fails basic disjoint n-amalgamation, since the first forbidden configuration has size n. However, K_n^k already fails disjoint $(k+1)$-amalgamation. Let $p_\varnothing$ be the type of a complete hypergraph on $(n-k-1)$ vertices. For $\varnothing \neq S \in \mathcal{P}^-([k+1])$, let p_S be the type of a complete hypergraph in the variables $x_\varnothing$ together with x_i for $i \in S$. This is fewer than n variables, so p_S is consistent. But for any k-tuple $\overline{y}$ from $\overline{x}_\varnothing$ and $x_0, \ldots, x_k$, $\overline{y}$ is contained in p_S for some $S \in \mathcal{P}^-([k+1])$, so p_S implies $R(\overline{y})$. Thus the p_S cannot be amalgamated consistently: any amalgam would imply the existence of a complete hypergraph on n vertices.

On the other hand, $\mathcal{K}_n^k$ has ℓ-DAP for all $2 \leq \ell \leq k$. Suppose $(p_S)_{S \subsetneq [\ell]}$ is a coherent family of types. Define $p_{[\ell]}$ to be the type extending $\bigcup_{S \subsetneq [\ell]} p_S$ by asserting that there are no edges other than those specified by the types p_S (this is essentially the "free amalgamation"). If $p_{[\ell]}$ is consistent, it is a solution to the disjoint ℓ-amalgamation problem. So assume for contradiction that $p_{[\ell]}$ is inconsistent. Then there are variables $y_1, \ldots, y_n$ such that $p_{[\ell]}$ says $y_1, \ldots, y_n$ are the vertices of a complete k-hypergraph. For each $i \in [\ell]$, let $S_i = [\ell] \setminus \{i\}$. Since p_{S_i} is complete and contained in $p_{[\ell]}$, p_{S_i} cannot mention all the variables y_j. So there is some y_{j_i} such that y_{j_i} is not in $\overline{x}_{S'}$ for any $S' \subseteq S_i$. Extend $y_{j_0}, \ldots, y_{j_{\ell-1}}$ arbitrarily to a k-tuple $\overline{y}$. Then $R(\overline{y}) \in p_{[\ell]}$. By definition of $p_{[\ell]}$, there is some $S \subsetneq [\ell]$ such that $R(\overline{y}) \in p_S$. But picking i such that $i \notin S$, $S \subseteq S_i$, and p_S does not mention the variable y_{j_i}, contradiction.

If we replace $\mathcal{P}^-([n])$ by another downwards closed family of subsets $\mathcal{F}$ in the definitions above, we call the amalgamation problem **partial**. It will be useful to observe that disjoint amalgamation gives solutions to partial amalgamation problems as well.

Lemma 7.5. *Suppose that T has (basic) k-DAP for all $2 \leq k \leq n$. Then every (basic) partial disjoint n-amalgamation problem has a solution.*

Proof. The same proof works in the basic case and the general case.

We are given a coherent $\mathcal{F}$-family of types $\{p_S \mid S \in \mathcal{F}\}$, with $\mathcal{F} \subseteq \mathcal{P}^-([n])$ downwards closed, and we seek a non-redundant type $p_{[n]}$ with $p_S \subseteq p_{[n]}$ for all $S \in \mathcal{F}$. Note that we can assume $\mathcal{F}$ is non-empty, otherwise any type $p_{[n]}$ is a solution.

We build a solution "from the bottom up". By induction on $0 \leq k \leq n$, I claim that we can extend this family to a coherent $\mathcal{F}_k$-family of types, where $\mathcal{F}_k = \mathcal{F} \cup \{S \subseteq [n] \mid |S| \leq k\}$. When $k = n$, we have a coherent $\mathcal{P}([n])$-family of types, as desired.

When $k = 0$, since $\mathcal{F}$ is non-empty and downwards closed, $\varnothing \in \mathcal{F}$, so $\mathcal{F}_0 = \mathcal{F}$. When $k = 1$, if there is any i such that $i \notin S$ for all $S \in \mathcal{F}$, then the original $\mathcal{F}$-family of types says nothing about variables $\overline{x}_{\{i\}}$. We choose any type $p_{\{i\}}(\overline{x}_\varnothing, x_i)$ in the single variable x_i (note that $\overline{x}_\varnothing$ is empty in the basic case).

Now given a coherent $\mathcal{F}_{k-1}$-family of types by induction, with $2 \leq k \leq n$, we wish to extend to a coherent $\mathcal{F}_k$-family of types. If there is any set $S \subseteq [n]$ with $|S| = k$ such that $S \notin \mathcal{F}_{k-1}$, then all proper subsets of S are in $\mathcal{F}_{k-1}$. Hence we have types $\{p_R \mid R \in \mathcal{P}^-(S)\}$ which form a coherent $\mathcal{P}^-(S)$-family. Using k-DAP, we can find a non-redundant type p_S in the variables $\overline{x}_S$ extending the types p_R. Doing this for all such S gives a coherent $\mathcal{F}_k$-family of types, as desired. □

We are largely interested in theories with n-DAP for all n, and in this case there is no difference between basic and full n-DAP.

Proposition 7.6. *T has n-DAP for all n iff T has basic n-DAP for all n.*

Proof. One direction is clear, since basic disjoint n-amalgamation is a special case of disjoint n-amalgamation.

In the other direction, note first that there is a solution to the disjoint n-amalgamation problem $\{p_S \mid S \in \mathcal{P}^-([n])\}$ if and only if the partial type

$$\{x \neq x' \mid x, x' \text{ distinct}\} \cup \bigcup_{S \in \mathcal{P}^-([n])} p_S$$

is consistent. Hence, by compactness, we can reduce to the case that each tuple of variables $\overline{x}_S$ is finite.

Let $N = \sum_{S \in \mathcal{P}^-([n])} |\overline{x}_S|$. Renumber the x variables $x_0, \ldots, x_{N-1}$. Now each type p_S determines a type in some subset of these variables. Closing downward under restriction to smaller sets of variables, we obtain a partial basic disjoint N-amalgamation problem. By Lemma 7.5 and basic disjoint N-amalgamation, this partial amalgamation problem has a solution, a type $p_{[N]}(x_0, \ldots, x_{N-1})$. This type is a solution to the original n-amalgamation problem. □

We now prove that if the theory of a Fraïssé limit has n-DAP for all n, then it is pseudofinite. The proof involves a probabilistic construction of a structure of size N for each N "from the bottom up". This is the same idea as in the proof of Lemma 7.5, but there we could fix an arbitrary k-type extending a given coherent family of l-types for $l < k$. Here we introduce randomness by choosing an extension uniformly at random. Note that the measures μ_N constructed in the proof are not necessarily the uniform measures on $\mathcal{K}(N)$, but we obtain a zero-one law for $(\mu_N)_{N \in \omega}$, which implies pseudofiniteness.

The probabilistic calculation is essentially the same as the one used in the proof of the zero-one law for graphs (Theorem 6.2). The key point is that the amalgamation properties allow us to make all choices as independently as possible: the quantifier-free types assigned to subsets A and B of $[N]$ are independent when conditioned on the quantifier-free type assigned to $A \cap B$. It is this independence which makes the calculation go through.

Theorem 7.7. *Let $\mathcal{K}$ be a finitary Fraïssé class such that $T^*_{\mathcal{K}} = \mathrm{Th}(M_{\mathcal{K}})$ has n-DAP for all n. Then every sentence in $T^*_{\mathcal{K}}$ has a finite model in $\mathcal{K}$.*

Proof. I will define a probability measure μ_N on the finite set $\mathcal{K}(N)$ for each $N \in \omega$ by describing a probabilistic construction of a structure $M_N \in \mathcal{K}(N)$.

We assign k-types to each subset of size k from $[N]$ by induction. Note that by QE, assigning a k-type to a tuple $i_0, \ldots, i_{k-1}$ from $[N]$ is the same as picking a structure $A \in \mathcal{K}(k)$. The k-type will be completely determined by $\theta_A(i_0, \ldots, i_{k-1})$.

When $k = 0$, there is no choice: by HP and JEP (or by completeness of $T^*_{\mathcal{K}}$), there is a unique empty structure in $\mathcal{K}(0)$, and a unique 0-type. When $k = 1$, for each $i \in [N]$, choose the 1-type of $\{i\}$ uniformly at random from $\mathcal{K}(1)$. Now proceed inductively: having assigned l-types to all subsets of size l with $l < k$, we wish to assign k-types. For each k-tuple $i_0, \ldots, i_{k-1}$ of distinct elements from $[N]$, let $P = \{p_S \mid S \in \mathcal{P}^-([k])\}$ be the collection of types assigned to all proper subtuples. Since $T^*_{\mathcal{K}}$ has k-DAP, there is a solution to this basic disjoint k-amalgamation problem. And there are finitely many, since they correspond to a subset $\mathcal{K}(k, P) \subseteq \mathcal{K}(k)$ of solutions to the disjoint k-amalgamation problem. We choose a solution from $\mathcal{K}(k, p)$ uniformly at random.

Continuing by induction to $k = N$, the resulting structure M_N is in $K(N)$. I claim that if φ is any axiom of $T^*_{\mathcal{K}}$ then $\lim_{N\to\infty} \mu_N([\varphi]) = 1$. This suffices, by Lemma 6.1.

Each universal axiom $\varphi \in T_{\mathcal{K}}$ has the form $\forall x_1, \ldots, x_k\, \psi(\overline{x})$, where ψ is quantifier-free and true on all k-tuples from structures in $\mathcal{K}$. Since all substructures the random structure M_N are in K, φ is always satisfied by M_N, and so $\mu_N([\varphi]) = 1$ for all N.

Now consider the one-point extension axiom $\varphi_{A,B}$: $\forall \overline{x}\, (\theta_A(\overline{x}) \to \exists y\, \theta_B(\overline{x}, y))$. Let $\overline{a}$ be a tuple of $|A|$-many distinct elements from $[N]$ and b any other element. Conditioning on the event that $M_N \models \theta_A(\overline{a})$, I claim there is a positive probability ε that $M_N \models \theta_B(\overline{a}, b)$.

Indeed, θ_B specifies the diagram of the tuple $\overline{a}b$ among those allowed by $\mathcal{K}$. There is a positive probability $(1/|\mathcal{K}(1)|)$ that the correct 1-type is assigned to b, and, given that the correct l-type has been assigned to all subtuples of $\overline{a}b$ involving b of length $l < k$, there is a positive probability $(1/|\mathcal{K}(k, P)|$ for the appropriate basic disjoint k-amalgamation problem $P)$ that the correct k-type is assigned to a given subtuple of length k involving b. Then ε is the product of all these probabilities for $1 \leq k \leq |B|$.

Moreover, for distinct elements b and b', the events that $\overline{a}b$ and $\overline{a}b'$ satisfy θ_B are conditionally independent, since the types of tuples involving elements from $\overline{a}$ and b but not b' are decided independently from those of tuples involving elements from $\overline{a}$ and b' but not b, conditioned on the type assigned to $\overline{a}$.

We finish with the same computation as in Theorem 6.2, by computing the probability that φ is *not* satisfied by M_N. Conditioned on the event that $M_N \models \theta_A(\overline{a})$, the probability that $M_N \not\models \exists y\, \theta_B(\overline{a}, y)$ is $(1-\varepsilon)^{N-|A|}$, since there are $N - |A|$ choices for the element b, each with independent probability $(1 - \varepsilon)$ of failing to satisfy θ_B. Removing the conditioning, the probability that $M_N \not\models \exists y\, (\theta_A(\overline{a}) \to \theta_B(\overline{a}, y))$ for any given $\overline{a}$ is at most $(1-\varepsilon)^{N-|A|}$, since the formula is vacuously satisfied when $\overline{a}$ does not satisfy θ_A. Finally, there are $N^{|A|}$ possible tuples $\overline{a}$, so the probability that $M_N \not\models \forall \overline{x}\, \exists y\, (\theta_A(\overline{x}) \to \theta_B(\overline{x}, y))$ is at most $N^{|A|}(1 - \varepsilon)^{N-|A|}$. Since $|A|$ is constant, the exponential decay dominates the polynomial growth, and $\lim_{N\to\infty} \mu_N([\neg\varphi]) = 0$, so $\lim_{N\to\infty} \mu_N([\varphi]) = 1$. □

Lecture 8: More on n-DAP and pseudofiniteness

Let's say a theory T has ω-DAP if it has n-DAP for all $n \geq 2$. By Morleyizing, we obtain the following consequence of Theorem 7.7.

Corollary 8.1. *Any $\aleph_0$-categorical theory T with ω-DAP is pseudofinite.*

Example 8.2. There are $\aleph_0$-categorical theories T which do not have ω-DAP but which admit $\aleph_0$-categorical expansions T' with ω-DAP. By Theorem 7.7, T' is pseudofinite. And since the reduct of a pseudofinite theory is pseudofinite, it follows that T is pseudofinite.

As a simple example, consider the theory of an equivalence relation with k infinite classes ($k \geq 2$). This theory fails to have 3-DAP just as in Example 7.3(1). But if we expand the language by adding k new unary relations $C_1, \ldots, C_k$ in such a way that each class is named by one of the C_i, the resulting theory has ω-DAP.

For a more interesting example, the random graph (which has n-DAP for all n, by Example 7.2) has a reduct to a 3-hypergraph in the language R', where the relation $R'(a, b, c)$ holds if and only if there are an *odd* number of the three possible edges between a, b, and c. This structure turns out to be homogeneous, hence the Fraïssé limit of its age, which is the class of all finite 3-hypergraphs with the property that on any four distinct vertices a, b, c, and d, there are an *even* number of the four possible 3-edges. Such a 3-hypergraph is (rather confusingly) called a "two-graph" in the literature. The class of finite two-graphs fails to have 4-DAP. Nevertheless, the theory of its Fraïssé limit is pseudofinite, as a reduct of the theory of the random graph.

This method of proving pseudofiniteness can be pushed a bit further.

Definition 8.3. A Fraïssé class $\mathcal{K}$ is *filtered* by a chain $\mathcal{K}_0 \subseteq \mathcal{K}_1 \subseteq \mathcal{K}_2 \subseteq \ldots$ if each $\mathcal{K}_n$ is a Fraïssé class, and $\bigcup_{n\in\omega} \mathcal{K}_n = \mathcal{K}$.

Proposition 8.4. *Let $\mathcal{K}$ be a finitary Fraïssé class filtered by $\{\mathcal{K}_n \mid n \in \omega\}$. For every sentence φ, $T^*_{\mathcal{K}} \models \varphi$ if and only if $T^*_{\mathcal{K}_n} \models \varphi$ for all sufficiently large n.*

Proof. In the forward direction, suppose $T^*_{\mathcal{K}} \models \psi$. By compactness, finitely many axioms $\varphi_1, \ldots, \varphi_k$ of $T^*_{\mathcal{K}}$ entail ψ. It suffices to show that each φ_i is entailed by $T^*_{\mathcal{K}_n}$ for all sufficiently large n, since then, for all sufficiently large n, $T^*_{\mathcal{K}_n} \models \bigwedge_{i=1}^{n} \varphi_i$, and $T^*_{\mathcal{K}_n} \models \psi$.

Since $\mathcal{K}_n \subseteq \mathcal{K}$ for all n, for each universal axiom $\psi \in T_{\mathcal{K}}$, $T_{\mathcal{K}_n} \models \psi$ for all n. Now let (A, B) be a one-point extension in $\mathcal{K}$ with corresponding axiom $\varphi_{A,B}$. For large enough n, the structures A and B are in $\mathcal{K}_n$, so (A, B) is also a one-point extension in $\mathcal{K}_n$, and $T^*_{\mathcal{K}_n} \models \varphi_{A,B}$.

For the backward direction, suppose $T^*_{\mathcal{K}_n} \models \psi$ for all sufficiently large n, and suppose for contradiction that $T^*_{\mathcal{K}} \not\models \psi$. Since $T^*_{\mathcal{K}}$ is complete, $T^*_{\mathcal{K}} \models \neg\psi$. By the forward direction, $T^*_{\mathcal{K}_n} \models \neg\psi$ for all sufficiently large n. This is a contradiction, since all $T^*_{\mathcal{K}_n}$ are consistent. $\square$

Corollary 8.5. *If a Fraïssé class $\mathcal{K}$ is filtered by $\{\mathcal{K}_n \mid n \in \omega\}$ and each generic theory $T_{\mathcal{K}_n}$ is pseudofinite, then the generic theory $T_{\mathcal{K}}$ is pseudofinite.*

Proof. If $T^*_{\mathcal{K}} \models \psi$, then $T^*_{\mathcal{K}_n} \models \psi$ for sufficiently large n by Proposition 8.4, and hence ψ has a finite model. □

As a consequence, if $\mathcal{K}$ is filtered by $\{\mathcal{K}_n \mid n \in \omega\}$, and each $T^*_{\mathcal{K}_n}$ admits an expansion to an $\aleph_0$-categorical theory with ω-DAP, then $T^*_{\mathcal{K}}$ is pseudofinite.

Example 8.6. Let $\mathcal{K}_E$ be the class of finite equivalence relations. $T^*_{\mathcal{K}_E}$ is the theory of an equivalence relation with infinitely many infinite classes. This theory has no $\aleph_0$-categorical expansion with ω-DAP. But we can filter $\mathcal{K}_E$ by letting $\mathcal{K}_n$ be the class of finite equivalence relations with $\leq n$ classes. Then $T^*_{\mathcal{K}_n}$ is the theory of an equivalence relation with n infinite classes. This theory has an $\aleph_0$-categorical expansion with ω-DAP by naming the classes (see Example 8.2 above). Corollary 8.5 implies that $T^*_{\mathcal{K}_E}$ is pseudofinite.

Of course, it is not hard to see directly that $T^*_{\mathcal{K}_E}$ is pseudofinite, since it is axiomatized by sentences asserting that there are $\geq n$ classes of size $\geq n$ for all n, and each of these axioms is true in a finite equivalence relation. We apply this method to a less trivial theory: the theory T^*_{feq} of generic parametrized equivalence relations.

Let $\mathcal{L}$ be the language with two unary predicates, O and P (for "objects" and "parameters"), and a ternary relation $E_x(y, z)$. Then $\mathcal{K}_{\mathrm{feq}}$ is the class of finite structures with the property that for all a in P, $E_a(y, z)$ is an equivalence relation on O.

Lemma 8.7. *$\mathcal{K}_{\mathrm{feq}}$ is a Fraïssé class with DAP. Moreover, given $f_1 \colon A \hookrightarrow B_1$ and $f_2 \colon A \hookrightarrow B_2$ and $n \in \omega$ such that for every $b \in P(B_1)$ or $P(B_2)$, E_b has at most n equivalence classes, there is a solution C to the amalgamation problem such that for all $c \in P(C)$, E_c has at most n equivalence classes.*

Proof. $\mathcal{K}_{\mathrm{feq}}$ clearly has HP. For DAP, suppose we have embeddings $f_1 \colon A \hookrightarrow B$ and $f_2 \colon A \hookrightarrow B_2$ of structures in $\mathcal{K}_{\mathrm{feq}}$, and assume that for every $b \in P(B_1)$ or $P(B_2)$,E_b has at most n equivalence classes. We specify a structure C with domain $A \cup (B_1 \setminus f_1(A)) \cup (B_2 \setminus f_2(A))$ into which B_1 and B_2 embed in the obvious way. For each parameter a in $P(C)$, we must specify an equivalence relation on $O(C)$ with at most n classes. If a is in $P(A)$, it already defines equivalence relations on B_1 and B_2. First, number the E_a-classes in A by $1, \ldots, \ell$. Then, if there are further unnumbered E_a-classes in B_1 and C_2, number them by $\ell + 1, \ldots, m_1$ and $\ell + 1, \ldots, m_2$ respectively. Note that $m_1, m_2 \leq n$. Now define E_a in $O(C)$ to have $\max(m_1, m_1)$ classes by merging the classes assigned the same number in the obvious way. The situation is even simpler if a is not in $P(A)$. Say without loss of generality it is in $P(B_1)$. Then we can extend E_a to $O(C)$ by adding all elements of $O(B_2 \setminus f_1(A))$ to a single existing E_a-class. JEP follows from AP by taking A to be the empty structure. □

We define T^*_{feq} to be the theory of the Fraïssé limit of $\mathcal{K}_{\mathrm{feq}}$. To show T^*_{feq} is pseudofinite, we filter the class $\mathcal{K}_{\mathrm{feq}}$ by the subclasses K_n in which each equivalence relation in the parameterized family has at most n classes. If we expand these classes by parameterized predicates naming each equivalence class, the resulting theory has ω-DAP.

Theorem 8.8. *T^*_{feq} is pseudofinite.*

Proof. For $n \in \omega$, let $\mathcal{K}_n$ be the subclass of $\mathcal{K}_{\text{feq}}$ consisting of those structures with the property that for all a in P, the equivalence relation E_a has at most n classes. It follows from Lemma 8.7 that each $\mathcal{K}_n$ is a Fraïssé class.

For any structure A in $\mathcal{K}_{\text{feq}}$, if $|O(A)| = N$, then for all $a \in P(A)$, the equivalence relation E_a has at most N classes, so $A \in K_N$. Hence $\mathcal{K}_{\text{feq}} = \bigcup_{n\in\omega} K_n$. So $\mathcal{K}_{\text{feq}}$ is a filtered Fraïssé class, and by Corollary 8.5, it suffices to show that each $T^*_{\mathcal{K}_n}$ is pseudofinite.

Let $\mathcal{L}'_n$ be the expanded language which includes, in addition to the relations O, P, and E, n binary relation symbols $C_1(x, y), \dots, C_n(x, y)$. Let $\mathcal{K}'_n$ be the class of finite L'_n-structures which are expansions of structures in $\mathcal{K}_n$ such that for all a of sort P, each of the E_a-classes is picked by one of the formulas $C_i(a, y)$. Then $\mathcal{K}'_n$ is a Fraïssé class by essentially the same argument as for $\mathcal{K}_n$. Let M_n and M'_n be the Fraïssé limits of $\mathcal{K}_n$ and $\mathcal{K}'_n$, respectively.

We certainly have $\mathcal{K}_n = \{A|_{\mathcal{L}} \mid A \in \mathcal{K}'_n\}$, since every structure in $\mathcal{K}_n$ can be expanded to one in $\mathcal{K}'_n$ by labeling the classes for each equivalence relation arbitrarily. Suppose now that (A, B) is a one-point extension in $\mathcal{K}_n$ and A' is an expansion of A to a structure in $\mathcal{K}'_n$. If the new element $b \in B$ is in $P(B)$, then it defines a new equivalence relation E_b on $O(A) = O(B)$, and we can expand B to B' in $\mathcal{K}'_n$ by labeling the E_b-classes arbitrarily. On the other hand, suppose b is in $O(B)$. Then for each parameter a, either b is an existing E_a-class labeled by $C_i(a, y)$, in which case we set $C_i(a, b)$, or b is in a new E_a-class, in which case we set $C_j(a, b)$ for some unused C_j. By a homework exercise, it follows that $M_n = M'_n|_{\mathcal{L}}$.

Finally, I claim that $T^*_{K'_n}$ has ω-DAP. It certainly has 2-DAP, since it is a Fraïssé class with the disjoint amalgamation property. The behavior of the ternary relation $E_x(y, z)$ is entirely determined by the behavior of the binary relations $C_i(x, y)$, and an L'_n-structure $(P(A), O(A))$ is in K'_n if and only if for every a in $P(A)$ and b in $O(a)$, $C_i(a, b)$ holds for exactly one i. So any inconsistency is already ruled out at the level of the 2-types. Since in a coherent $\mathcal{P}^-([n])$-family of types for $n \geq 3$, every pair of variables is contained in one of the types, we conclude that there are no inconsistencies, and every disjoint n-amalgamation problem has a solution.

So $T_{K'_n}$ has disjoint n-amalgamation for all n, and hence it and its reduct T_{K_n} are pseudofinite by Theorem 7.7. □

Remark 8.9. A natural question is whether T^*_{feq} is, in fact, the almost-sure theory for the class $\mathcal{K}_{\text{feq}}$ for the uniform measures. Consider the sentence

$$\forall x\, \forall x'\, \forall y\, \forall y'\, ((P(x) \wedge P(x') \wedge O(y) \wedge O(y') \wedge x \neq x') \to \exists z\, (E_x(y, z) \wedge E_{x'}(y', z))),$$

which expresses that any two equivalence classes for distinct equivalence relations intersect. This sentence is in T^*_{feq}, since it is implied by the relevant one-point extension axioms. But it has limit probability 0 with respect to the uniform measures.

The fact from enumerative combinatorics is that the expected number of equivalence classes in an equivalence relation on a set of size n, chosen uniformly, grows asymptotically as $\frac{n}{\log(n)}(1+o(1))$. Thus, most of the equivalence relations E_a have equivalence classes which are much smaller (with average size approximately $\log(n)$) than the number of classes, and the probability that every E_a-class is large enough to intersect every E_b-class nontrivially for all distinct a and b converges to 0.

It's a fact that if an $\aleph_0$-categorical theory T has 3-DAP, it is simple with trivial forking (supersimple of U-rank 1 and trivial acl). Simplicity is a property that lives strictly between stability and NSOP_3.

In fact, Shelah defined three more properties between simplicity and NSOP_3, called SOP_1, SOP_2, and TP_1. It is now known that all three of these properties are equivalent (in the sense that if a theory has any one of them, it has all three). I won't define these properties (or simplicity) here, but note that SOP_1 and SOP_2 are defined in a completely different way than SOP_n for $n \geq 3$. We have

$$\text{stable} \Rightarrow \text{simple} \Rightarrow \mathrm{NTP}_1 = \mathrm{NSOP}_{1,2} \Rightarrow \mathrm{NSOP}_3 \Rightarrow \mathrm{NSOP}_4 \Rightarrow \cdots \Rightarrow \mathrm{NSOP}.$$

It is an open problem whether $\mathrm{NTP}_1 = \mathrm{NSOP}_3$.

Recall that we know examples of pseudofinite $\aleph_0$-categorical theories that are stable (like the infinite pure set or the equivalence relation with infinitely many infinite classes), and simple (like the random graph). The theory T^*_{feq} is NTP_1 and not simple. It was the first known example of a non-simple pseudofinite $\aleph_0$-categorical theory.

We saw earlier that if T is pseudofinite, then it is NSOP. As far as I know, for *every* $\aleph_0$-categorical theory T with NSOP and SOP_3, pseudofiniteness of T is an open problem. I do not know of any such theory for which the answer is known (in either direction!).

Lecture 9: Weak amalgamation and genericity

We return now to the general setting of $\mathcal{K}$-limits for an age $\mathcal{K}$.

Given an age $\mathcal{K}$ and a property P of $\mathcal{K}$-limits, consider the following game $G(\mathcal{K}, P)$ for two players. Player I picks a structure $A_0 \in \mathcal{K}$. Then Player II picks a structure $A_1 \in \mathcal{K}$ containing A_0 as a substructure. Then Player I picks a structure $A_2 \in \mathcal{K}$ containing A_1 as a substructure. The play continues in this way, with the players taking turns constructing a chain

$$A_0 \subseteq A_1 \subseteq A_2 \subseteq \ldots$$

After ω turns, we consider the $\mathcal{K}$-limit $A_\omega = \bigcup_{i\in\omega} A_i$ (see Lemma 1.6). Player II wins if A_ω satisfies P. Otherwise, Player I wins.

Definition 9.1. We say the property P is **generic** (for $\mathcal{K}$) if Player II has a winning strategy in $G(\mathcal{K}, P)$. We say a $\mathcal{K}$-limit M is **generic** if the property $A_\omega \cong M$ is generic.

Formally, a strategy is a function mapping a partial play of the game

$$A_0 \subseteq \cdots \subseteq A_n$$

to a structure $A_{n+1} \in \mathcal{K}$ containing A_n as a substructure.

For those who know some basic descriptive set theory, it is possible to define a topological space $\mathrm{Lim}_\mathcal{K}$ of $\mathcal{K}$-limits with domain ω in such a way that (under some mild assumptions on $\mathcal{K}$) the game $G(\mathcal{K}, P)$ is essentially the Banach–Mazur game in $\mathrm{Lim}_\mathcal{K}$. Then a property P is generic if and only if the set of $\mathcal{K}$-limits satisfying P is comeager in $\mathrm{Lim}_\mathcal{K}$, and a $\mathcal{K}$-limit M is generic if and only if the isomorphism class of M is comeager in $\mathrm{Lim}_\mathcal{K}$. But the embedding game $G(\mathcal{K}, P)$ has less technical baggage, so we will stick to this perspective for now.

Example 9.2. If $\mathcal{K}$ is an age, then the property $\mathrm{Age}(A_\omega) = \mathcal{K}$ is generic. Player II can follow the strategy in the proof of Theorem 1.3, using JEP each turn to embed the isomorphism representative B_n from $\mathcal{K}$ into A_{2n+1}, and hence into A_ω.

Example 9.3. If $\mathcal{K}$ is a Fraïssé class, then the Fraïssé limit $M_\mathcal{K}$ is generic. Player II can follow the strategy in the proof of Theorem 2.5. After Player I produces a structure A_{2n}, Player II can look at all f.g. substructures B of A_{2n} and all embeddings $B \hookrightarrow C$ with C an isomorphism representative from $\mathcal{K}$ and adds these to the list of tasks. By completing a task on teach turn, Player II can force A_ω to be rich (and hence isomorphic to $M_\mathcal{K}$) no matter how Player I plays.

Example 9.4. It is possible for an age $\mathcal{K}$ to have a generic limit even when no Fraïssé limit exists (i.e., when it fails to have AP). For example, the infinitely branching countable tree is a generic limit for the class $\mathcal{K}_{\text{forests}}$. Indeed, on turn $2n + 1$, given a forest A_{2n}, Player II can add vertices to make the forest connect, and then add further neighbors to every vertex in A_{2n} to ensure that these vertices have degree $\geq n$. The direct limit of connected forests (trees) is connected (a tree), and every vertex in the limit has infinite degree. So A_ω is the infinitely branching countable tree. We saw in Examples 2.4 and 2.8 that $\mathcal{K}_{\text{forests}}$ fails to have AP and has no Fraïssé limit.

Lemma 9.5. *If Player II has a winning strategy in the game $G(\mathcal{K}, P)$, then Player I has a winning strategy in the game $G(\mathcal{K}, \neg P)$. The converse is true when P is isomorphism invariant and $\mathcal{K}$ is an age.*

Proof. Suppose Player II has a winning strategy σ for $G(\mathcal{K}, P)$. We define a winning strategy σ' for Player I in $G(\mathcal{K}, \neg P)$. Player I picks an arbitrary structure $A_0 \in \mathcal{K}$ to be his first move, so writing ε for the empty play, $\sigma'(\varepsilon) = A_0$. Then we define $\sigma'(A_0 \subseteq A_1 \subseteq \cdots \subseteq A_{2n+1}) = \sigma(A_1 \subseteq \cdots \subseteq A_{2n+1})$. That is, Player I "steals" Player II's strategy, and plays as if he were the second player.

After a run in which Player I uses σ', we have $A_\omega = \bigcup_{n\in\omega} A_n = \bigcup_{n\geq 1} A_n$, and the chain $A_1 \subseteq A_2 \subseteq \dots$ is a run in which Player II uses σ, so A_ω satisfies P. Thus Player I wins $G(\mathcal{K}, \neg P)$.

Conversely, suppose Player I has a winning strategy σ for $G(\mathcal{K}, \neg P)$, and assume $\mathcal{K}$ is an age and P is isomorphism invariant. We define a winning strategy σ' for Player II in $G(\mathcal{K}, P)$. Let $A_* = \sigma(\varepsilon)$, Player I's first move according to σ. Given any structure $A_0 \in \mathcal{K}$, by JEP we can pick $A_1 \in \mathcal{K}$ with $A_0 \subseteq A_1$ and an embedding $f\colon A_* \hookrightarrow A_1$, and we can define $\sigma'(A_0) = A_1$. On all subsequent turns, given $A_0 \subseteq \cdots \subseteq A_{2n}$, for $n > 1$ define A'_i to be a structure isomorphic to A_i but with the elements of $f(A_*)$ replaced by the elements of A_*. Then if $\sigma(A_* \subseteq A'_2 \subseteq \cdots \subseteq A'_{2n}) = A'_{2n+1}$, define A_{2n+1} to be a structure isomorphic to A'_{2n+1} but with the elements of A_* replaced by the elements of $f(A_*)$. Let $\sigma'(A_0 \subseteq \cdots \subseteq A_{2n}) = A_{2n+1}$.

After a run in which Player II uses σ', since $A'_n \cong A_n$ for all $n \geq 2$, we have $A_\omega = \bigcup_{n\in\omega} A_n = \bigcup_{n\geq 2} A_n \cong A_* \cup \bigcup_{n\geq 2} A'_n$, and the chain $A_* \subseteq A'_2 \subseteq A'_3 \subseteq \dots$ is a run in which Player I uses σ, so (by isomorphism invariance) A_ω satisfies P and Player II wins $G(\mathcal{K}, P)$. $\square$

A property P is **determined** if one of the players has a winning strategy in $G(\mathcal{K}, P)$. A non-determined property will correspond to a subset of $\mathrm{Lim}_\mathcal{K}$ without the Baire property, so it is necessarily somewhat exotic. Under mild assumptions on $\mathcal{K}$, every sentence of $\mathcal{L}_{\omega_1,\omega}$ defines a Borel subset of $\mathrm{Lim}_\mathcal{K}$ and hence is determined. Since truth of sentences is isomorphism-invariant, the set of generic sentences is a complete theory, even in first-order logic or even in $\mathcal{L}_{\omega_1,\omega}$.

Corollary 9.6. *Let $\mathcal{K}$ be an age. For every determined isomorphism-invariant property P, either P or $\neg P$ is generic for $\mathcal{K}$.*

Proof. If Player II has a winning strategy in $G(\mathcal{K}, P)$, then P is generic. Otherwise, since P is determined, Player I has a winning strategy in $G(\mathcal{K}, P)$. By Lemma 9.5, since $\mathcal{K}$ has JEP and P is isomorphism invariant, Player II has a winning strategy in $G(\mathcal{K}, \neg P)$. So $\neg P$ is generic. $\square$

Theorem 9.7. *If $\mathcal{K}$ has a generic limit, it is unique up to isomorphism.*

Proof. Suppose M and N are both generic $\mathcal{K}$-limits. Let σ be a winning strategy for Player II in $G(\mathcal{K}, \cong M)$, and let τ be a winning strategy for Player II in $G(\mathcal{K}, \cong N)$. By Lemma 9.5, Player I has a winning strategy τ' in $G(\mathcal{K}, \ncong N)$. Let $A_0 \subseteq A_1 \subseteq \dots$ be a run in which Player I uses τ' and Player II uses σ. Then this play is winning for Player I in $G(\mathcal{K}, \ncong N)$ and winning for Player II in $G(\mathcal{K}, \cong M)$ so $M \cong A_\omega \cong N$. $\square$

Our goal now is to characterize when $\mathcal{K}$ has a generic limit.

Definition 9.8. Let $\mathcal{K}$ be an age and $A \in \mathcal{K}$.

- An **extension** of A is a pair (B, f), where $B \in \mathcal{K}$ and $f : A \hookrightarrow B$ is an embedding. We define a preorder on embeddings by $(B, f) \leq (C, g)$ if and only if there exists $h : B \hookrightarrow C$ such that $h \circ f = g$. We could just as well define a category of extensions (in category-theoretic language, this is the slice category under A), but the preorder structure is all we need.
- Two extensions (C_1, g_1) and (C_2, g_2) of A are **compatible** if there exists (D, h) with $(C_1, g_1) \leq (D, h)$ and $(C_2, g_2) \leq (D, h)$.
- An extension (B, f) is a **guide** for A if for all $(B, f) \leq (C_1, g_1)$ and $(B, f) \leq (C_2, g_2)$, (C_1, g_1) and (C_2, g_2) are compatible.
- A is an **amalgamation base** if (A, id_A) is a guide for A.

Intuitively, a guide (B, f) has "enough information" about A to ensure amalgamation over A. If A is an amalgamation base, then A already holds all the information about itself.

Definition 9.9. $\mathcal{K}$ has the **weak amalgamation property (WAP)** if every $A \in \mathcal{K}$ has a guide (B, f).

Unpacking the definition, we have the following: For all $A \in \mathcal{K}$, there exists $f : A \hookrightarrow B$ in $\mathcal{K}$ such that for all $g_1 : B \hookrightarrow C_1$ and $g_2 : B \hookrightarrow C_2$ in $\mathcal{K}$ there exist $h_1 : C_1 \hookrightarrow D$ and $h_2 : C_2 \hookrightarrow D$ in $\mathcal{K}$ such that $h_1 \circ g_1 \circ f = h_2 \circ g_2 \circ f$.

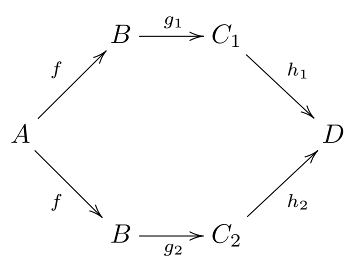

It is possible that in the amalgam D, we have $g_1(h_1(B)) \neq g_2(h_2(B))$, i.e., the diagram may not commute over B.

Example 9.10. The class $\mathcal{K}_{\text{forests}}$ has WAP. For any forest A, there is an embedding $f : A \hookrightarrow B$, where B is connected (a tree). Then (B, f) is a guide for A. In fact, any tree B is an amalgamation base in $\mathcal{K}_{\text{forests}}$.

Definition 9.11. $\mathcal{K}$ has the **cofinal amalgamation property (CAP)** if every $A \in \mathcal{K}$ embeds in some amalgamation base.

Note that $\mathcal{K}$ has AP if and only if every $A \in \mathcal{K}$ is an amalgamation base. And if $f : A \hookrightarrow B$ and B is an amalgamation base, then (B, f) is a guide for A. So AP implies CAP implies WAP.

We have seen that $\mathcal{K}_{\text{forests}}$ has CAP but not AP. I will now give an example with WAP but not CAP.

Example 9.12. Define a ternary relation R on $\mathbb{Q}$ by $R(x, y, z)$ if and only if $x < z$ and $y < z$ and $x \neq y$. Intuitively, $R(x, y, z)$ means "x,y,z are distinct and z is greatest." Let $\mathcal{K} = \mathrm{Age}(\mathbb{Q}, R)$. $\mathcal{K}$ has WAP: Given $A \in \mathcal{K}$, we can pick $f\colon A \hookrightarrow (\mathbb{Q}, R)$ and pick some $b \in \mathbb{Q}$ with $b < f(a)$ for all $a \in A$. Then the embedding of A in $B = f(A) \cup \{b\}$ is a guide for A. But $\mathcal{K}$ has no amalgamation base B with $|B| \geq 2$, so $\mathcal{K}$ does not have CAP.

Lecture 10: Weak richness and weak homogeneity

We now introduce the weakenings of $\mathcal{K}$-richness and $\mathcal{K}$-homogeneity that correspond to the weak amalgamation property.

Definition 10.1. Let M be a $\mathcal{K}$-limit.

- M is **weakly $\mathcal{K}$-rich** if for all $A \subseteq_{\text{f.g.}} M$ there exists $A \subseteq B \subseteq_{\text{f.g.}} M$ such that for any $f\colon B \hookrightarrow C$ with $C \in \mathcal{K}$, there exists $g\colon C \hookrightarrow M$ such that $(g \circ f)|_A = \mathrm{id}_A$. We call B an **embedded guide** for A.
- M is **weakly $\mathcal{K}$-homogeneous** if $\mathrm{Age}(M) = \mathcal{K}$ and for all $A \subseteq_{\text{f.g.}} M$ there exists $A \subseteq B \subseteq_{\text{f.g.}} M$ such that for any embedding $g\colon B \hookrightarrow M$, there exists $\sigma \in \mathrm{Aut}(M)$ such that $\sigma|_A = g|_A$.

Note that the terminology "embedded guide" is different from the usage of "guide" in the definition of WAP, though they will turn out to be the same in the proofs.

A bit of history: The notion of weak $\mathcal{K}$-homogeneity (with $\mathcal{K} = \mathrm{Age}(M)$) was introduced by Pabion under the name "prehomogeneity" in 1972. Example 9.12 above appeared the 1972 paper of Pabion, attributed to Pouzet. In 1994, Pouzet and Roux proved the equivalence between prehomogeneity and genericity.

Independently, in the 1990s there began to be a great deal of interest in generic automorphisms of homogeneous structures. Truss showed in 1992 that CAP (for the class of structures in a Fraïssé class $\mathcal{K}$ expanded by a partial automoprhism) was sufficient to obtain a generic automorphism. The theorem (originally due to Pouzet and Roux) that WAP characterizes the existence of a generic limit was again obtained independently by Ivanov in 1999 (who called it the "almost amalgamation property") and by Kechris and Rosdendal in 2007 (under the name WAP).

Lemma 10.2. *If $\mathcal{K}$ is an age and M is a weakly rich $\mathcal{K}$-limit, then* $\mathrm{Age}(M) = \mathcal{K}$.

Proof. Since M is a $\mathcal{K}$-limit, $\mathrm{Age}(M) \subseteq \mathcal{K}$. Let $A \in \mathcal{K}$, and let $E = \langle \varnothing \rangle \subseteq_{\text{f.g.}} M$. Let $E \subseteq B \subseteq_{\text{f.g.}} M$ be a witness to weak $\mathcal{K}$-richness for M. By JEP for $\mathcal{K}$, there is some $C \in \mathcal{K}$ and embeddings $f\colon B \hookrightarrow C$ and $g\colon A \hookrightarrow C$. By weak $\mathcal{K}$-richness, there is an embedding $h\colon C \hookrightarrow M$ such that $(h \circ f)|_E = \mathrm{id}_E$. Then $h \circ g\colon A \hookrightarrow M$ shows $A \in \mathrm{Age}(M)$. ☐

Theorem 10.3. *Let M and N be weakly rich $\mathcal{K}$-limits. Suppose we have $A \subseteq B \subseteq_{\text{f.g.}} M$ and $A' \subseteq B' \subseteq_{\text{f.g.}} N$, and $f\colon A \cong A'$ and $g\colon B \cong B'$ are* *isomorphisms with $f \subseteq g$. Assume B is an embedded guide for A. Then there is an isomorphism $\varphi\colon M \cong N$ such that $\varphi|_A = f$.*

Proof. Just as in Theorem 1.13, we enumerate $M = (m_i)_{i\in\omega}$ and $N = (n_i)_{i\in\omega}$ and build an isomorphism by a back-and-forth argument. This time, we define a sequence of isomorphisms $f_i\colon A_i \cong A'_i$ and $g_i\colon B_i \cong B'_i$, where $A_i \subseteq B_i \subseteq_{\text{f.g.}} M$ and $A'_i \subseteq B'_i \subseteq_{\text{f.g.}} N$. We require that $f_i \subseteq g_i$, $f \subseteq f_i \subseteq f_{i+1}$, $n_i \in A'_{2i+1}$, and $m_i \in A_{2i+2}$ for all $i \in \omega$. Note that we do not require $g_i \subseteq g_{i+1}$. When i is even, we assume B_i is an embedded guide for A_i, and when i is odd, we assume B'_i is an embedded guide for A'_i.

At the end of the construction, we have $\bigcup_{i\in\omega} A_i = M$, $\bigcup_{i\in\omega} A'_i = N$, and $\bigcup_{i\in\omega} f_i\colon M \cong N$ extending f.

Set $A_0 = A$, $B_0 = B$, $A'_0 = A'$, $B'_0 = B'$, $f_0 = f$, and $g_0 = g$.

At odd stage $2i + 1$, we are given $f_{2i}\colon A_{2i} \cong A'_{2i}$ and $g_{2i}\colon B_{2i} \cong B'_{2i}$, with B_{2i} an embedded guide for A_{2i}. Define $A'_{2i+1} = \langle B'_{2i} \cup \{n_i\}\rangle \subseteq_{\text{f.g.}} N$, and let $A'_{2i+1} \subseteq B'_{2i+1} \subseteq_{\text{f.g.}} N$ be an embedded guide for A'_{2i+1}. Since N is a $\mathcal{K}$-limit, $A'_{2i+1}, B'_{2i+1} \in \mathcal{K}$.

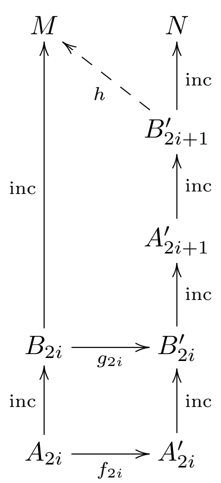

Since B_{2i} is an embedded guide for A_{2i}, and $g_{2i}\colon B_{2i} \hookrightarrow B'_{2i+1}$, there exists $h\colon B'_{2i+1} \hookrightarrow M$ with $(h \circ g_{2i})|_{A_{2i}} = \mathrm{id}_{A_{2i}}$. Let $A_{2i+1} = h(A'_{2i+1})$ and $B_{2i+1} = h(B'_{2i+1})$. Then $f_{2i+1} = (h|_{A'_{2i+1}})^{-1}$ and $g_{2i+1} = (h|_{B'_{2i+1}})^{-1}$ are isomorphisms $A_{2i+1} \cong A'_{2i+1}$ and $B_{2i+1} \cong B'_{2i+1}$ with $f_{2i+1} \subseteq g_{2i+1}$. Since $f_{2i} = g_{2i}|_{A_{2i}}$, $f_{2i} \subseteq f_{2i+1}$. And B'_{2i+1} is an embedded guide for A'_{2i+1} by construction.

The construction is similar at even stage $2i + 2$. □

Corollary 10.4. *Let M be a $\mathcal{K}$-limit. Then M is weakly rich if and only if M is weakly homogeneous. Moreover if M and N are both weakly rich/homogeneous, then $M \cong N$.*

Proof. First, assume M is weakly homogeneous. Let $A \subseteq_{\text{f.g.}} M$. Then there exists $A \subseteq B \subseteq_{\text{f.g.}} M$ witnessing weak homogeneity. I claim that B is an embedded guide for A. Let $f\colon B \hookrightarrow C$ with $C \in \mathcal{K}$. Since $\mathrm{Age}(M) = \mathcal{K}$

(this is part of the definition of weakly $\mathcal{K}$-homogeneous), there is an embedding $g\colon C \hookrightarrow M$, but we may not have $(g \circ f)|_A = \mathrm{id}_A$. Since $(g \circ f)|_B\colon B \hookrightarrow M$, by weak homogeneity, there exists $\sigma \in \mathrm{Aut}(M)$ such that $\sigma|_A = (g \circ f)|_A$. Let $g' = \sigma^{-1}|_{g(C)} \circ g\colon C \hookrightarrow M$. Then $(g' \circ f)|_A = (\sigma^{-1} \circ g \circ f)|_A = \mathrm{id}_A$.

The rest is a consequence of Theorem 10.3. Suppose M is weakly rich. By Lemma 10.2, $\mathrm{Age}(M) = \mathcal{K}$. Let $A \subseteq_{\text{f.g.}} M$, and let B be an embedded guide for A. Given $g\colon B \hookrightarrow M$, we have $g\colon B \cong g(B)$ and $g|_A\colon A \cong g(A)$, so by Theorem 10.3, there is an isomorphism $\sigma \in \mathrm{Aut}(M)$ with $\sigma|_A = g|_A$.

Having proved that weakly rich and weakly homogeneous are equivalent properties for $\mathcal{K}$-limits, let M and N be weakly rich. Pick some $A \subseteq B \subseteq_{\text{f.g.}} M$ such that B is an embedded guide for A. Since $\mathrm{Age}(N) = \mathcal{K}$, there is an embedding $g\colon B \hookrightarrow N$. Let $B' = g(B)$, so $g\colon B \cong B'$, and let $A' = g(A)$ and $f = g|_A\colon A \cong A'$. By Theorem 10.3, $M \cong N$. □

It follows from the proof of Corollary 10.4 that the witnesses to weak homogeneity are exactly the embedded guides witnessing weak richness.

Theorem 10.5. *Let $\mathcal{K}$ be an age with WAP. The following property of $\mathcal{K}$-limits is generic: M is weakly rich, and whenever $A \subseteq B \subseteq_{\text{f.g.}} A_\omega$ and the inclusion $A \hookrightarrow B$ is a guide, then B is an embedded guide for A.*

Proof. We define a winning strategy for Player II in the game $G(\mathcal{K}, \star)$, where $\star$ is the property in the statement of the theorem Just as in Example 9.3, we turn the task structure in the proof of Theorem 2.5 into a strategy, but modified to replace richness with weak richness.

We fix an enumeration $(D_k)_{k\in\mathcal{K}}$ of the structures in $\mathcal{K}$ up to isomorphism.

Given a partial play of the game $C_0 \subseteq \cdots \subseteq C_{2n}$, where Player I has just played C_{2n}, Player II enumerates all triples (A_j, B_j, g_j) where $A_j \subseteq B_j \subseteq C_{2n}$, the inclusion $A_j \hookrightarrow B_j$ is a guide, and $g_j\colon B_j \hookrightarrow D_{k_j}$ is an embedding into one of our isomorphism representatives. Player II adds these triples as tasks (n, j) to an array of tasks.

Player II then picks task $t(n) = (i, j)$ from the list. This is a triple (A_j, B_j, g_j), where $A_j \subseteq B_j \subseteq C_{2i} \subseteq C_{2n}$, the inclusion $A_j \hookrightarrow B_j$ is a guide, and $g_j\colon B_j \hookrightarrow D_{k_j}$. Since $A_j \hookrightarrow B_j$ is a guide, (C_{2n}, inc) and $g_j|_{A_j}$ are compatible. Pick some E_{2n+1} with embeddings $h'\colon C_{2n} \hookrightarrow E_{2n+1}$ and $h_{2n+1}\colon D_j \hookrightarrow E_{2n+1}$. Finally, let $h''\colon E_{2n+1} \hookrightarrow C_{2n+1}$ be a guide for E_{2n+1}. We may assume that h' and h'' are inclusions, so $C_{2n} \subseteq E_{2n+1} \subseteq C_{2n+1}$.

After a run in which Player II follows this strategy, we construct a $\mathcal{K}$-limit $C_\omega = \bigcup_{n\in\omega} C_n$. I claim that C_ω satisfies $\star$.

First, we show that if $A \subseteq B \subseteq_{\text{f.g.}} C_\omega$, and the inclusion $A \subseteq B$ is a guide, then B is an embedded guide for A. Since B is f.g., there is some $i \in \omega$ such that $B \subseteq C_{2i}$. Let $g\colon B \hookrightarrow D$ with $D \in \mathcal{K}$. We may assume that $D = D_k$ is one of our isomorphism representatives. Then the triple (A, B, g) was added to the list of tasks at some stage, and completed at some stage n. So there is an embedding $h_{2n+1}\colon D_k \hookrightarrow C_{2n+1}$ such that $h_{2n+1} \circ g$ is equal to the inclusion of A in C_{2n+1}, as desired.

It remains to show that every $A \subseteq_{\text{f.g.}} M$ has an embedded guide. Since A is f.g., there is some $i \in \omega$ such that $A \subseteq C_{2i} \subseteq E_{2i+1} \subseteq C_{2i+1} \subseteq C_\omega$. Since the inclusion $E_{2i+1} \subseteq C_{2i+1}$ is a guide for E_{2i+1}, the inclusion $A \subseteq C_{2i+1}$ is a guide for A. From what we just showed above, C_{2i+1} is an embedded guide for A. $\square$

Theorem 10.6. *Suppose $\mathcal{K}$ is an age which does not have WAP. Let M be a $\mathcal{K}$-limit. The property that A_ω embeds in M is not generic.*

Proof. We show Player I has a winning strategy in $G(\mathcal{K}, \exists f\colon A_\omega \hookrightarrow M)$.

Since $\mathcal{K}$ does not have WAP, there exists $A \in \mathcal{K}$ with no guide. Player I plays $A_0 = A$ on his first move. Let $(f_i)_{i\in\omega}$ enumerate the embeddings $A \hookrightarrow M$. On turn $2n$ with $n > 0$, Player II has just played A_{2n-1} with $A \subseteq A_{2n-1}$. Since the inclusion $\text{inc}\colon A \hookrightarrow A_{2n-1}$ is not a guide for A, there exist $g_i\colon A_{2n-1} \hookrightarrow C_i$ with $i = 1, 2$ such that $g_1|_A$ and $g_2|_A$ are not compatible. We may assume each g_i is an inclusion.

If there is an embedding $h\colon C_1 \hookrightarrow M$ with $(h \circ g_1)|_A = f_n$, Player I plays $A_{2n} = C_2$. Otherwise, Player I plays $A_{2n} = C_1$.

Now after a run on which Player I follows this strategy, suppose for contradiction that there exists $f\colon A_\omega \hookrightarrow M$. Then $f|_A\colon A \hookrightarrow M$ is equal to f_n for some $n \in \omega$. Let $f' = f|_{A_{2n}}$. If $A_{2n} = C_1$, then $(f' \circ g_1)|_A = f_A = f_n$, so $A_{2n} = C_2$, contradiction. Thus $A_{2n} = C_2$, and we also have $h\colon C_1 \hookrightarrow M$ such that $(h \circ g_1)|_A = f_n$. Letting $D = \langle h(C_1) \cup f'(C_2)\rangle \subseteq_{\text{f.g.}} M$, $h\colon C_1 \hookrightarrow D$ and $f'\colon C_2 \hookrightarrow D$ show that C_1 and C_2 are compatible over A, contradiction. $\square$

Theorem 10.7. *Let $\mathcal{K}$ be an age. The following are equivalent.*

(1) $\mathcal{K}$ has WAP.

(2) Weak richness is a generic property of $\mathcal{K}$-limits.

(3) There exists a weakly rich $\mathcal{K}$-limit.

(4) There exists a generic $\mathcal{K}$-limit.

Proof. (1)⇒(2): This is Theorem 10.5.

(2)⇒(3) is clear.

(3)⇒(1): Let M be a weakly rich $\mathcal{K}$-limit. Let $A \in \mathcal{K}$. Then there is an embedding $f\colon A \hookrightarrow M$. Let $f(A) \subseteq B \subseteq_{\text{f.g.}} M$ be an embedded guide for $f(A)$. I claim that $f\colon A \hookrightarrow B$ is a guide for A. Given $g_1\colon B \hookrightarrow C_1$ and $g_2\colon B \hookrightarrow C_2$ with $C_1, C_2 \in \mathcal{K}$, by weak $\mathcal{K}$-richness there exist $h_1\colon C_1 \hookrightarrow M$ and $h_2\colon C_2 \hookrightarrow M$ such that $(h_i \circ g_i)|_{f(A)} = \text{id}_{f(A)}$ for $i = 1, 2$. Let $D = \langle C_1 \cup C_2\rangle \subseteq_{\text{f.g.}} M$. Then $D \in \mathcal{K}$, $h_i\colon C_i \to D$ for $i = 1, 2$, and $h_1 \circ g_1 \circ f = h_2 \circ g_2 \circ f$.

(2)⇒(4): Suppose weak richness is a generic property. Let M be a weakly rich $\mathcal{K}$-limit. Player II has a winning strategy in the game $G(\mathcal{K}, \text{weakly rich})$, and by Corollary 10.4, any weakly rich $\mathcal{K}$-limit is isomorphic to M, so Player II has a winning strategy in the game $G(\mathcal{K}, \cong M)$. Thus M is generic.

(4)⇒(1): Suppose M is a generic $\mathcal{K}$-limit. Then the property $\cong M$ is generic, so the weaker property $\exists f\colon A_\omega \hookrightarrow M$ is generic. By Theorem 10.6, $\mathcal{K}$ has WAP. $\square$

Remark 10.8. It follows from the proof of Theorem 10.7 that if M is a generic $\mathcal{K}$-limit, then for $A \subseteq B \subseteq_{\text{f.g.}} M$, the inclusion $A \subseteq B$ is a guide for A if and only if B is an embedded guide for A. Indeed, we proved in Theorem 10.5 that it is generic that every guide is an embedded guide, hence this is true in the generic $\mathcal{K}$-limit. Conversely, since weak richness is generic, M is weakly rich, and we proved in the (3)⇒(1) direction of Theorem 10.7 that embedded guides are guides in weakly rich $\mathcal{K}$-limits.

Lecture 11: Universal and generic limits, generic automorphisms

Corollary 11.1. *If $\mathcal{K}$ has a universal limit, then $\mathcal{K}$ has a generic limit.*

Proof. Suppose $\mathcal{K}$ has a universal limit M, and assume for contradiction that $\mathcal{K}$ has no generic limit. By Theorem 10.7, $\mathcal{K}$ does not have WAP. By Theorem 10.6, the property that A_ω embeds in M is not generic. Thus there exists some $\mathcal{K}$-limit which does not embed in M, contradicting universality. □

By Corollary 11.1, whenever $\mathcal{K}$ has a universal limit, it must have a generic limit. Sometimes the generic limit is itself universal – for example, this happens when $\mathcal{K}$ has AP, since the Fraïssé limit is both generic and universal. But the generic limit need not be universal in general! We have the following implications, all of which are strict:

$\mathcal{K}$ has a Fraïssé limit	$\Longrightarrow$	$\mathcal{K}$ has a universal generic limit	$\Longrightarrow$	$\mathcal{K}$ has a universal limit	$\Longrightarrow$	$\mathcal{K}$ has a generic limit

Example 11.2. Recall that a graph is **acyclic** if it has no cycles. A graph is **linear** if every vertex has degree at most 2. Normally we consider graphs in the language $\mathcal{L} = \{R\}$, but a **colored graph** is a graph in the language $\text{Ł}_P = \{R, P\}$ augmented by an additional unary predicate (which is interpreted arbitrarily).

- The class of finite graphs has a Fraïssé limit (the random graph), as does the class of colored graphs (this was a problem on Homework 1).
- The class of finite ayclic graphs (forests) has no Fraïssé limits, but it has a universal generic limit, namely the countable infinitely branching tree. Genericity was Example 9.4, and universality was Example 2.4. The same is true for finite acyclic colored graphs: the universal generic limit is the countable infinitely branching tree colored so that each vertex has infinitely many neighbors satisfying P and infinitely many satisfying $\neg P$. In both cases, a guide for A is any finite tree containing A.

- The class of finite linear acyclic graphs has a universal limit (and hence a generic limit), but the generic is not universal. The generic limit is a single infinite chain (the unique infinite connected linear acyclic graph), since Player II has a winning strategy in the game $G(\mathcal{K}, \text{connected})$. A guide for A is a connected finite chain containing A. The universal limit is the disjoint union of countably many infinite chains. Note that the universal limit does not embed in the generic limit.

- The class of finite linear graphs (not necessarily acyclic) has the same behavior as the previous example. The property of connectedness is no longer generic, but the property that every vertex is contained in a finite cycle is. A guide for A is a graph containing A which is a disjoint union of cycles. The generic limit is the disjoint union of countably many cycles of each finite size ≥ 3. The universal limit is the same, but with infinitely many chains as well.

- The class of finite colored linear graphs has a generic limit but no universal limit. The generic limit is the disjoint union of countably many copies of each finite colored cycle. Again, a guide for A is a graph containing A which is a disjoint union of cycles. To see that there is no universal limit, note that we can find $2^{\aleph_0}$-many colored infinite chains, none of which embeds in any other. A countable limit contains at most countably many infinite chains, and hence cannot be universal.

- Finally, the class of finite colored linear acyclic graphs has no generic limit. We show that WAP fails. Let A be a graph with a single vertex v. Suppose for contradiction that $f\colon A \hookrightarrow B$ is a guide for A. By extending further, we may assume B is connected and that it is a chain with v as its midpoint. Let C_1 extend B by adding vertices satisfying P to both ends of the chain. Let C_2 extend B by adding vertices satisfying $\neg P$ to both ends of the chain. Then C_1 and C_2 are not compatible extensions of A.

In the case of finitary ages, we can characterize the behaviors above in terms of the spaces of existential types.

Definition 11.3. Let T be a theory. An **existential type** is a set of existential formulas. An existential type is **maximal** if it is maximal under containment among existential types consistent with T. Let $S_n^{\exists}(T)$ be the spaces of maximal existential types in n free variables. We topologize $S_n^{\exists}(T)$ by taking the basic open sets to be $[\varphi(x_1, \ldots, x_n)] = \{p \in S_n^{\exists}(T) \mid \varphi \in p\}$ where $\varphi(x_1, \ldots, x_n)$ is an existential formula. An existential type p is **isolated** if there is an existential formula φ such that $[\varphi] = \{p\}$.

Fact 11.4. Suppose $\mathcal{K}$ is a finitary age. We work relative to $T_{\mathcal{K}}$, the universal theory of $\mathcal{K}$-limits.

(1) $\mathcal{K}$ has a generic limlt if and only if the isolated types are dense in $S_n^{\exists}(T_{\mathcal{K}})$ for all n.

(2) $\mathcal{K}$ has a universal limit if and only if $S_n^{\exists}(T_{\mathcal{K}})$ is countable for all n.

(3) $\mathcal{K}$ has a universal-generic limit if and only if every type is isolated in $S_n^{\exists}(T_{\mathcal{K}})$ for all n.

(4) $\mathcal{K}$ has a Fraïssé limit if and only if every type is isolated by a quantifier-free formula in $S_n^{\exists}(T_{\mathcal{K}})$ for all n.

Note the analogy with countable models of a countable complete theory T: case (1) corresponds to the case when T has a prime model, case (2) corresponds to the case when T has a countable saturated model, and case (3) corresponds to the case when T is $\aleph_0$-categorical. Then case (4) corresponds to the case when T is $\aleph_0$-categorical with quantifier elimination. Note that $S_n^{\exists}(T_{\mathcal{K}})$ is not necessarily compact, so it is possible to have infinitely many types in $S_n^{\exists}(T_{\mathcal{K}})$, each of which is isolated.

When $S_n^{\exists}(T_{\mathcal{K}})$ is countable, then $T_{\mathcal{K}}$ in fact admits a countable existentially-saturated existentially closed model, which is a universal $\mathcal{K}$-limit, and which is unique up to isomorphism. This gives a canonical universal $\mathcal{K}$-limit.

It is natural to wonder whether cases (2) and (3) in the Fact can be characterized by "finitary diagrammatic properties" like AP and WAP. We interpret the following theorem (which is as-yet unpublished joint work with Aristotelis Panagiotopoulos) as saying the answer is no.

Theorem 11.5. *Let $\mathcal{L} = \{R\}$, where R is a binary relation. In a suitable (Polish) space of ages in the language $\mathcal{L}$, the sets $\{\mathcal{K} \mid \mathcal{K}$ has a universal limit$\}$ and $\{\mathcal{K} \mid \mathcal{K}$ has a universal-generic limit$\}$ are complete coanalytic (Π_1^1-complete).*

Many interesting examples of generic limits come from generic automorphisms.

Let $\mathcal{K}$ be a Fraïssé class $\mathcal{L}$. Let $\mathcal{L}_p = \mathcal{L} \cup \{p\}$, where p is a new binary relation symbol. Let $\mathcal{K}_p$ be the class of (A, p) where $A \in \mathcal{K}$ and p is the graph of a partial isomorphism, i.e., an $\mathcal{L}$-isomorphism $\varphi\colon B \cong C$ where B and C are substructures of A. The class $\mathcal{K}_p$ is almost never a Fraïssé class.

Definition 11.6. We say that the Fraïssé limit $M_{\mathcal{K}}$ admits a **generic automorphism** if $\mathcal{K}_p$ has a generic limit (M, p), where $M \cong M_{\mathcal{K}}$ and p is the graph of an automorphism of M.

The existence of a generic automorphism of $M_{\mathcal{K}}$ is equivalent to the existence of $\sigma \in \mathrm{Aut}(M_{\mathcal{K}})$ such that the conjugacy class $\{\tau\sigma\tau^{-1} \mid \tau \in \mathrm{Aut}(M_{\mathcal{K}})\}$ is comeager in the topological group $\mathrm{Aut}(M_{\mathcal{K}})$.

Let us first consider generic automorphisms of $(\mathbb{Q}, <)$.

For any linear order $(L, <)$ and an automorphism $\sigma \in \mathrm{Aut}(L)$, we define an equivalence relation $\sim_\sigma$ by $a \sim_\sigma b$ if and only if there exists $n \in \mathbb{Z}$ such that $\sigma^n(a) \leq b \leq \sigma^{n+1}(a)$ or $\sigma^{n+1}(a) \leq b \leq \sigma^n(a)$. In other words, b is contained in the convex hull of the orbit of a. Note that $\sim_\sigma$ partitions L into convex equivalence classes. There are three types of classes:

(1) If $\sigma(a) = a$, then $[a]_{\sim_\sigma} = \{a\}$.

(2) If $a < \sigma(a)$, then for all n, $\sigma^n(a) < \sigma^{n+1}(a)$. If $b \sim_\sigma a$, with $\sigma^n(a) < b < \sigma^{n+1}(a)$, then $b < \sigma^{n+1}(a) < \sigma(b)$, so $b < \sigma(b)$. Thus σ is increasing on all of $[a]_{\sim_\sigma}$.

(3) If $\sigma(a) < a$, a similar argument shows that σ is decreasing on all of $[a]_{\sim_\sigma}$.

A generic automorphism of $(\mathbb{Q}, <)$ is one in which all three types of equivalence classes described above are dense in the order on the equivalence classes (i.e., between any two equivalence classes, classes of all three types appear). In fact, letting p be the graph of a generic automorphism, $(\mathbb{Q}, <, p)$ is a universal limit of $\mathcal{K}_{\mathrm{LO},p}$.

The random graph also has a generic automorphism, but the picture is rather different.

Definition 11.7. A Fraïssé class $\mathcal{K}$ has the **extension property for partial automorphisms (EPPA)** if every $A \in \mathcal{K}$ embeds in some $B \in \mathcal{K}$ such that every partial isomorphism of A extends to an automorphism of B.

Note that K_{LO} does not have EPPA. Every finite linear order is rigid, so non-trivial partial isomorphisms of A cannot extend to automorphisms of any finite B.

The following theorem is due to Hrushovski. I'll present a simpler proof due to Herwig and Lascar.

Theorem 11.8. *The class $\mathcal{K}_{\mathrm{G}}$ of finite graphs has EPPA.*

Proof. First, for a finite set X and $n \in \omega$, we define $\Gamma(X, n)$ to be the graph whose vertices are subsets of X of size n, such that there is an edge between Y and Z if and only if $Y \cap Z \neq \varnothing$. Note that if σ is a permutation of X, then the map $\sigma^*\colon \Gamma(X, n) \to \Gamma(X, n)$ by $Y \mapsto \sigma(Y)$ is an automorphism of $\Gamma(X, n)$.

Now let $A = (V, R)$ be a finite graph. Let $n \in \omega$ such that $n \geq 2$ and $n \geq \deg(v)$ for all $v \in A$. We turn A into a multi-graph with self-loops: For each $v \in A$ with $\deg(v) < n$, add self-loops to v until $\deg(v) = n$. The resulting multi-graph is $A' = (V, R')$. Now let $B = \Gamma(R', n)$.

Define $f\colon A \hookrightarrow B$ by $f(v) = \{e \in R' \mid v \in e\}$. This is an embedding. Indeed, if $v \neq w$, since $\deg_{R'}(v) = \deg_{R'}(w) = n > 2$, and there is at most one edge between v and w, $f(v) \neq v(w)$. And we have $\{v, w\} \in R$ if and only if $\{v, w\} \in f(v) \cap f(w)$ if and only if there is an edge between $f(v)$ and $f(w)$ in B.

Now let C and D be subgraphs of A and $\varphi\colon C \to D$ an isomorphism. We use φ to define a permutation of R'. For each edge $e = \{v, w\}$ in C, we map e to $\varphi(e) = \{\varphi(v), \varphi(w)\}$, which is an edge in D. For each $v \in C$, let $R'_{C,v}$ be the set of edges in R' containing v and no other element of C, and similarly for D. Note that the sets $R'_{C,v}$ are pairwise disjoint, and $|R'_{C,v}| = |R'_{D,\varphi(v)}|$, since $\deg_{R'}(v) = \deg_{R'}(\varphi(v)) = n$. So we can pick a bijection $R'_{C,v} \to R'_{D,\varphi(v)}$ arbitrarily for each v. Finally, we extend the partial permutation of R' constructed so far to an arbitrary permutation $\sigma\colon R' \to R'$, and thus obtain an automorphism σ^* of $\Gamma(R', n)$.

It remains to check that σ^* extends φ, in the sense that for all $v \in C$, $\sigma^*(f(v)) = f(\varphi(v))$. Indeed, for each edge $e \in f(v)$, if $e \in C$, then since $v \in e$, $\varphi(v) \in \varphi(e)$, so $\sigma(e) \in f(\varphi(v))$. And if $e \notin C$, then $e \in R'_{C,v}$, so $\sigma(e) \in R'_{D,\varphi(v)}$, and $\sigma(e) \in f(\varphi(v))$. It follows that $\sigma^*(f(v)) = f(\varphi(v))$, as desired. □

Definition 11.9. Given a Fraïssé class $\mathcal{K}$ and $n \in \omega$, let $\mathcal{K}_{p,n}$ be the class of all $(A, p_1, \ldots, p_n)$, where $A \in \mathcal{K}$ and each p_i is the graph of a partial isomorphism. We say $\mathcal{K}$ has **ample generic automorphisms** (or **ample generics**) if for every $n \in \omega$, $\mathcal{K}_{p,n}$ has a generic limit $(M, p_1, \ldots, p_n)$, where $M \cong M_{\mathcal{K}}$ and each p_i is the graph of an automorphism of M.

Corollary 11.10. *$\mathcal{K}_G$ has ample generics.*

Proof. With $\mathcal{K} = \mathcal{K}_G$, it suffices to show that $\mathcal{K}_{p,n}$ has WAP. Let $(A, R, p_1, \ldots, p_n) \in \mathcal{K}_{p,n}$. Let $f\colon (A, R) \hookrightarrow (B, R)$ witness EPPA. For each partial isomorphism p_i, pick some extension of it to an automorphism $\sigma_i \in \mathrm{Aut}(B)$. I claim that $(B, R, \sigma_1, \ldots, \sigma_n)$ is an amalgamation base. Indeed, given two extensions C_1 and C_2 of B, we can amalgamate C_1 and C_2 freely over B. □

Note that the generic limit $(G, R, p_1, \ldots, p_n)$ can be built as a chain of finite EPPA witnesses, in which the partial isomorphisms coded by the p_i are in fact total. It follows that every orbit for each generic automorphism is finite. Thus in this case the generic limit $(G, R, p_1, \ldots, p_n)$ is not universal, since there exist automorphisms of countable graphs with infinite orbits.