Limits of sparse hypergraphs

Riley Thornton

rthornto@andrew.cmu.edu
Carnegie Mellon University

Caltech logic seminar, November 2024

The plan

Ergodic theory has a useful notion of ultraproduct for automorphisms of measure spaces. We're going to steal it to do some combinatorics.

There's a nice class of structures on measure spaces generated by automorphisms (including all the finite structures!) We'll define ultraproducts for these.

This ends up recovering compactness (and more!) for the topology of local-global convergence.

I. Ultraproducts of maps

We view measure spaces through a functional analytic duality:

$$
\begin{aligned}
(X, \mu) &\leftrightsquigarrow L^2(X, \mu) \\
A &\leftrightsquigarrow \mathbf{1}_A \\
A \cap B &\leftrightsquigarrow \mathbf{1}_A \mathbf{1}_B \\
\mu(A) &\leftrightsquigarrow \langle \mathbf{1}_A, \mathbf{1} \rangle \\
T[A] &\leftrightsquigarrow T \cdot \mathbf{1}_A
\end{aligned}
$$

Proposition

If (X,μ) *is standard , then* $T \in \mathrm{Aut}(X,\mu)$ *is fixpoint free if and only if*

$$(\exists f_1, f_2, f_3)\ f_i^2 = f_i,\ f_i f_j = (T \cdot f_i) f_i = \mathbf{0},\ f_1 + f_2 + f_3 = \mathbf{1}.$$

"Standard" here means $L^2(X,\mu)$ is separable.

Proposition

If G *is a graph on a standard Borel space with degrees at most* d*, then* G *has a Borel* $d+1$*-coloring.*

Proposition

If (X, μ) *is standard , then* $T \in \mathrm{Aut}(X, \mu)$ *is fixpoint free if and only if*

$$(\exists f_1, f_2, f_3)\ f_i^2 = f_i,\ f_i f_j = (T \cdot f_i) f_i = \mathbf{0},\ f_1 + f_2 + f_3 = \mathbf{1}.$$

Proof.

Of course if T is not fixpoint free, we can't find characteristic functions like this that cover X.

If T is fixpoint free, then the Scheier graph T is degree at most 2, and we can 3-color it. We can take f_i to be the characteristic function of the i^{th} color class. □

Proposition

To any sequence of probability spaces (X_i, μ_i) we can associate an ultraproduct $(X, \mu) = \lim_{\mathcal{U}}(X_i, \mu_i)$ so that

1. *Any bounded sequence $\langle f_i : i \in I \rangle$ with $f_i \in L^2(X_i, \mu_i)$ has an ultralimit $\lim_{\mathcal{U}} f_i = f \in L^2(X, \mu)$*
2. *Any element of $L^2(X, \mu)$ is an ultralimit*
3. *Any sequence $T_i \in \mathrm{Aut}(X_i, \mu_i)$ has a limit $\lim_{\mathcal{U}} T_i = T$*
4. *All operations commute with limits, e.g.*

$$T(f + 3g)h = \lim_{\mathcal{U}} T_i(f_i + 3g_i)h_i.$$

In the special case where all T_i's are the same, we call this an ultrapower of T, $T^{\mathcal{U}}$.

Proposition

If T is an automorphism of a standard space, the ultrapower $T^{\mathcal{U}}$ is fixpoint free if T is fixpoint free.

Proof.

If T is fixpoint free, then T has a 3 independent characterstic functions that cover X, so the same is true of $T^{\mathcal{U}}$. □

On the other hand, ergodicity and mixing do not pass to the ultraproduct (these are omitting types properties).

II. Probability measure preserving structures.

We'll deal with a class of structures on measure spaces that generalize the following:

1. Any finite structure with normalized counting measure
2. The Schreier graph of a pmp group action
3. The Bernoulli enrichment of a transitive countable structure

The important points are that we can uniformly enumerate the (hyper)edges a vertex sees and that we have some kind of handshake lemma or unimodularity condition.

Definition

A **probability measure preserving (pmp) equivalence relation** is an equivalence relation E on a probability space (X, μ) so that $E = \bigcup_{i\in\omega} f_i$ where each f_i is a measure-preserving involution.

For a relation E,

$$(E)^n := \{(x_1, ..., x_n) : x_1 E x_2 E ... E x_n\}.$$

So, if $A \subseteq (E)^n$, any $x \in X$ can enumerate the A-edges its contained in.

Definition

Define a measure $\tilde{\mu}$ on $(E)^n$ by

$$\tilde{\mu}_i(A) = \int_{x \in X} |\pi_i^{-1}[x] \cap A| \, d\mu.$$

We have a measure theoretic handshake lemma:

Proposition

For any i, j, $\tilde{\mu}_i(A) = \tilde{\mu}_j(A)$.

So, for instance, there is no measurable directed graph in $(E)^2$ with outdegree 2 and indegree 1 everywhere.

Definition

A **pmp relation** is a Borel subset of $(E)^n$ for some n, E.

Equivalently, a pmp relation is a Borel relation generated by pmp involutions.

Equivalently, a pmp relation is a Borel relation satisfying the measurable handshake lemma.

We want to take ultraproducts of pmp structures. The basic idea is that a pmp structure is generated by a list of measure preserving involutions, so we can take the ultraproduct of these:

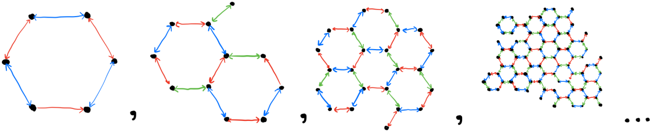

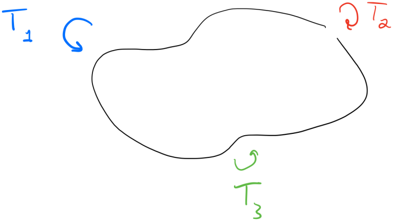

The problem is that there isn't a natural way to enumerate these involutions:

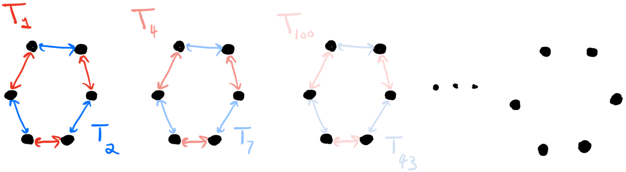

This isn't a problem if there are only finitely many T_i's.

Definition

For $T_1, T_2, T_3, \ldots \in \mathrm{Aut}(X, \mu)$, $\vec{j} \in \omega^k$ and $x \in X$, write

$$T_{\vec{j}}(x) = (T_{j_1}(x), \ldots, T_{j_k}(x)).$$

An n-marking of pmp relation $R \subseteq (E)^k$ on (X, μ) is a list of n automorphisms $T_1, \ldots, T_n$ so that, for all x,

$$\pi_1^{-1}[x] \cap R \subseteq \{ T_{\vec{j}}(x) : \vec{j} \in [n]^k \}.$$

By the same coloring proposition as before,

Proposition

For any d, u, there is an n so that every $R \subseteq (E)^u$ of degree at most d has an n marking.

Definition

For a sequence of pmp relations R_i, $T^i_1, \ldots, T^i_n$ an n-marking of R_i, and an ultrafilter $\mathcal{U}$, define the ultraproduct $\lim_{\mathcal{U}} R_i$ by

$$T_j = \lim_{\mathcal{U}} T^i_j$$

$$T_{\vec{j}}(x) \in \lim_{\mathcal{U}} R :\Leftrightarrow \lim_{\mathcal{U}} \mathbf{1}_{R_i}(T^i_{\vec{j}}(x)) = 1$$

This doesn't really depend on the marking.

Proposition

For a pmp graph $\mathcal{G}$ and an ultrafilter $\mathcal{U}$, the indepence ratio of $\mathcal{G}$ is the same as the independence ratio of $\mathcal{G}^{\mathcal{U}}$

Proposition

For a pmp graph $\mathcal{G}$ and an ultrafilter $\mathcal{U}$, the indepence ratio of $\mathcal{G}$ is the same as the independence ratio of $\mathcal{G}^{\mathcal{U}}$

Proof.

A function $f \in L^2(X, \mu)$ is the characteristic function of an independent set iff $f^2 = f$ and

$$(\forall x, T_i)\ f(x) = 0 \text{ or } f(T_i(x)) = 0 \text{ or } R(x, T_i(x)) = 0$$

and the measure of f is $\mathrm{tr}(f) = \langle f, \mathbf{1} \rangle$.

□

Proposition

For a pmp graph $\mathcal{G}$ and an ultrafilter $\mathcal{U}$, the indepence ratio of $\mathcal{G}$ is the same as the independence ratio of $\mathcal{G}^{\mathcal{U}}$

Proof.

A function $f \in L^2(X, \mu)$ is the characteristic function of an independent set iff $f^2 = f$ and

$$\prod_i f(T_i^{-1} \cdot f)(\mathbf{1}_R \circ (\mathrm{id}, T_i)) = 0$$

and the measure of f is $\mathrm{tr}(f) = \langle f, \mathbf{1} \rangle$.

□

Proposition

For a pmp graph $\mathcal{G}$ and an ultrafilter $\mathcal{U}$, the indepence ratio of $\mathcal{G}$ is the same as the independence ratio of $\mathcal{G}^{\mathcal{U}}$

Proof.

A function $f \in L^2(X, \mu)$ is the characteristic function of an independent set iff $f^2 = f$ and

$$\prod_i f(T_i^{-1} \cdot f)(\mathbf{1}_R \circ (\mathrm{id}, T_i)) = 0$$

and the measure of f is $\mathrm{tr}(f) = \langle f, \mathbf{1} \rangle$.
All of this algebra commutes with ultralimits, so if f witness $\alpha(\mathcal{G}) \geq r$, then $\lim_{\mathcal{U}} f$ witnesses $\alpha(\mathcal{G}^{\mathcal{U}}) \geq r$, and if $\lim_{\mathcal{U}} f_i$ witnesses $\alpha(\mathcal{G}^{\mathcal{U}}) \geq r$, then for some i, f_i witnesses $\alpha(\mathcal{G}) \geq r$. $\square$

More generally:

Proposition

Fix a rooted k-labelled finite structure (H, o, g). Then, for any $f : X \to [k]$

$$\mathbb{P}_\mu((B_r(x), x, f) \cong (H, o, g)) = \lim_{\mathcal{U}} \mathbb{P}_{\mu_i}((B_r(x), x, f_i) \cong (H, o, g))$$

So independence ratio, matching ratio, girth, etc all commute with ultraproducts.

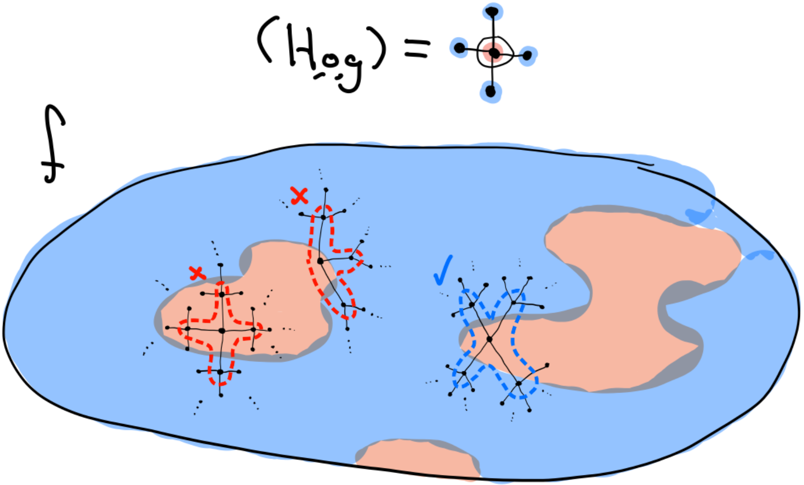
(Hog) =
f

III. Local-global convergence

Definition

For a pmp structures $\mathcal{H}$, f a measurable labelling of $\mathcal{H}$, the r-**local statistics** of f is the measure $\mu_r(f, \mathcal{H})$ on isomorphism types of rooted labelled finite structures in the same signature given by

$$\mu_r(f, \mathcal{H})([H, o, g]) = \mathbb{P}((B_r(x), x, f) \cong (H, o, g)).$$

Let $\mathcal{L}_{r,k}(\mathcal{H})$ be the closure of the set of r-local statistics of k-labellings of $\mathcal{H}$ in the total variation metric.
Two structures $\mathcal{H}, \mathcal{G}$ are **local-global equivalent** if $\mathcal{L}_{rk}(\mathcal{G}) = \mathcal{L}_{rk}(\mathcal{H})$ for all r, k.

We can put a topology on structures. The definition is clunky, but two structures are close if they can approximate most of the same local statistics.

Definition

Let $\mathbb{L}(u, d)$ be the set of relations on standard probability spaces with arity at most u and degree at most d. The **local-global topology**, $\mathbb{S}_{u,d}$ on $\mathbb{L}(u, d)$ is the coarsest topology making each $\mathcal{L}_{rk}$ a continuous map to the space of compact sets of measures with the Hausdorff metric.

This is not a Hausdorff space, but if we quotient out by local-global equivalence it becomes one.

Some examples:

Proposition

If Γ *is a countable amenable group, then the Schreier graph of any two free pmp actions of* Γ *are local-global equivalent, and they're the local-global limit of the induced subgraph on any Følner sequence.*

Proposition

If $\mathcal{H} = \lim_{\leftarrow} \mathcal{H}_i$ *is an inverse limit, then* $\mathcal{H}$ *is a local-global limit of* $\langle \mathcal{H}_i : i \in \omega \rangle$

Proposition (c.f. Hatami–Lovasz–Szegedy)

For any u, d, $\mathbb{S}_{u,d}$ is a compact psuedometric space.

Proof.

Since we've fixed a degree bound, there finitely many possible structures of a given size. So, the total variation distance on possible local statistics is compact. This means $\mathbb{S}$ is a precompact pseudometric space.

For compactness, given any sequence of structures $\mathcal{H}_i \in \mathbb{S}$, some subsequence has local statistics converging to those of $\lim_{\mathcal{U}} \mathcal{H}_i$.

And, by a Lowenheim–Skolem argument, we can find an equivalent structure on a standard space. □

We can also use some standard ultraproduct facts to get new theorems about local–global convergence and equivalence.

Proposition

If ν is a local statistic of $\mathcal{H}$, then there is $\mathcal{H}'$ equivalent to $\mathcal{H}$ and a labeling f of $\mathcal{H}$ so that $\mu_r(f) = \nu$.

Proposition

Any property of structures defined by an existential property in $L^2(X, \mu)$ (equipped with markings of the structure) is continuous with respect to local-global convergence

For instance, the spectral radius of a pmp graph is local-global convergent.

IV. An application

Definition

For a finite relational structure $\mathcal{D}$, an **instance** of $\mathrm{CSP}(\mathcal{D})$ is a structure in the same signature and a **solution** to instance $\mathcal{X}$ is a homomorphism from $\mathcal{X}$ to $\mathcal{D}$. (I'll sometimes say an instance of $\mathcal{D}$)

For example, an instance of the complete graph K_n is another graph and a solution is an n-coloring.

The width-1 structures are the ones where a certain simple constraint propagation algorithm (called arc-consistency) solves $\mathrm{CSP}(\mathcal{D})$.

Theorem

For a finite relational structure $\mathcal{D}$, the following are equivalent:

1. *$\mathcal{D}$ is width-1*
2. *Any standard measurable instance of $\mathcal{D}$ with a solution has a measurable solution*
3. *For any bounded degree sequence of instances $\mathcal{X}_i$, if $\lim_i \mathbb{P}(B_r(x) \cong \mathcal{H}) = 0$ for all unsolvable $\mathcal{H}$ and $r \in \mathbb{N}$, then there maps $f_i : \mathcal{X}_i \to \mathcal{D}$ so that f_i satisfies a proportion of constraints tending to 1*

Sketch.
($1 \to 2$) : The arc-constistency algorithm can be carried out measurably.

($2 \to 3$) : Suppose there were a sequence of counterexamples to (3). Then any limit point would be a counterexample to (2).

($3 \to 1$) : This boils down to building large girth hypergraphs with large chromatic number. A random regular hypergraph will do. □

Thanks!