Pseudo-random Graphs and Structures with Good Enough Measures

Henry Towsner[1]

University of Pennsylvania

June 14, 2013

[1]Supported by the US National Science Foundation DMS-1001528

Definition

We define the *triangles* in E by:

$$\triangle(E) = \{(x, y, z) \in V^3 \mid (x, y), (y, z), (z, x) \in E\}.$$

Theorem (Triangle Removal Lemma, Szemerédi and Rusza)

For every $\epsilon > 0$ there is a $\delta > 0$ such that for any finite V and any graph E on V, if $\mu^3(\triangle(E)) < \delta$ then

> *There is an $S \subseteq E$ with $\mu^2(S) < \epsilon$ such that $\triangle(E \setminus S) = \emptyset$ (it is possible to remove every triangle by removing a small number of edges).*

$\mu^2(S) = \frac{|S|}{|V|^2}$ and $\mu^3(\Delta(E)) = \frac{|\Delta(E)|}{|V|^3}$

Combinatorial results about density generally have analogs in which hold in *sparse* but *pseudorandom* structures.

Combinatorial results about density generally have analogs in which hold in *sparse* but *pseudorandom* structures.

For example, results about graphs (V, E) where V is a large finite set and $\mu^2(E) = \frac{|E|}{|V|^2} > \epsilon$ have analogs where

- $E \subseteq \Gamma \subseteq [V]^2$,
- $\mu^2(\Gamma)$ is very small—say $|\Gamma| \approx \gamma |V|^c$ with $c < 2$, but
- $\mu^2_\Gamma(E) = \frac{|E|}{|\Gamma|} > \epsilon$, and
- Γ is sufficiently pseudorandom.

Combinatorial results about density generally have analogs in which hold in *sparse* but *pseudorandom* structures.

For example, results about graphs (V, E) where V is a large finite set and $\mu^2(E) = \frac{|E|}{|V|^2} > \epsilon$ have analogs where

- $E \subseteq \Gamma \subseteq [V]^2$,
- $\mu^2(\Gamma)$ is very small—say $|\Gamma| \approx \gamma |V|^c$ with $c < 2$, but
- $\mu^2_\Gamma(E) = \frac{|E|}{|\Gamma|} > \epsilon$, and
- There is a set of graphs $\mathcal{P}$ such that $\Gamma \in \mathcal{P}$ and when Γ is chosen randomly,

$$\lim_{|V| \to \infty} \mathbb{P}(\Gamma \in \mathcal{P}) = 1.$$

Theorem (Kohayakawa, Rödl, Schacht, and Skokan, 2010)

For every $\epsilon > 0$ there is a $\delta > 0$ such that for any finite V, any sufficiently pseudorandom $\Gamma \subseteq [V]^2$, and any $E \subseteq \Gamma$, if $\mu^3_\Gamma(\triangle(E)) < \delta$ then

> *There is an $S \subseteq E$ with $\mu^2_\Gamma(S) < \epsilon$ such that $\triangle(E \setminus S) = \emptyset$ (it is possible to remove every triangle by removing a small number of edges).*

Here $\mu^3_\Gamma(\triangle(E)) = \frac{|\triangle(E)|}{|\triangle(\Gamma)|}$. If Γ is chosen randomly so $|\Gamma| \geq C|V|^{3/2}$ then, as $|V| \to \infty$, the probability that the theorem applies to Γ approaches 1.

Generalizations of triangle removal include:

Theorem (Graph Removal Lemma)

Let H be a fixed finite graph. For every $\epsilon > 0$ there is a $\delta > 0$ such that for any finite V and any graph E on V, if the measure of the set of copies of H (in $\mu^{|H|}$) is $< \delta$ then there is an $S \subseteq E$ with $\mu^2(S) < \epsilon$ such that there are no copies of H in $E \setminus S$.

Generalizations of triangle removal include:

Theorem (Graph Removal Lemma)

Let H be a fixed finite graph. For every $\epsilon > 0$ there is a $\delta > 0$ such that for any finite V and any graph E on V, if the measure of the set of copies of H (in $\mu^{|H|}$) is $< \delta$ then there is an $S \subseteq E$ with $\mu^2(S) < \epsilon$ such that there are no copies of H in $E \setminus S$.

Theorem (Hypergraph Removal Lemma)

Let H be any fixed finite k-regular hypergraph. For every $\epsilon > 0$ there is a $\delta > 0$ such that for any finite V and any k-regular hypergraph E on V, if the measure of the set of copies of H (in $\mu^{|H|}$) is $< \delta$ then there is an $S \subseteq E$ with $\mu^k(S) < \epsilon$ such that there are no copies of H in $E \setminus S$.

Theorem (T. 2012, Conlon-Fox-Zhou 2012)

Let H be a fixed finite graph. For every $\epsilon > 0$ there is a $\delta > 0$ such that for any finite V, any sufficiently pseudorandom $\Gamma \subseteq [V]^2$, and any $E \subseteq \Gamma$, if the measure of the set of copies of H (in $\mu_\Gamma^{|H|}$) is $< \delta$ then there is an $S \subseteq E$ with $\mu_\Gamma^2(S) < \epsilon$ such that there are no copies of H in $E \setminus S$.

Theorem (T. 2012)

Let H be any fixed finite k-regular hypergraph. For every $\epsilon > 0$ there is a $\delta > 0$ such that for any finite V, any sufficiently pseudorandom $\Gamma \subseteq [V]^k$, and any $E \subseteq K$, if the measure of the set of copies of H (in $\mu_\Gamma^{|H|}$) is $< \delta$ then there is an $S \subseteq E$ with $\mu_\Gamma^k(S) < \epsilon$ such that there are no copies of H in $E \setminus S$.

One approach to theorems like this is to use ultraproducts: suppose that for every n, (V_n, E_n) is a graph with $\mu^2(E_n) \geq \epsilon > 0$ and $|V_n| \to \infty$.

Then in the ultraproduct $(V, E) = \prod_{\mathfrak{U}}(V_n, E_n)$, the measures $\mu^1(S) = \prod_{\mathfrak{U}} \frac{|S_n|}{|V_n|}$ and $\mu^2(T) = \prod_{\mathfrak{U}} \frac{|T_n|}{|V_n|^2}$ extend to actual σ-additive measures and $\mu^2(E) \geq \epsilon > 0$.

We can prove theorems about (V_n, E_n) for large n by using measure theoretic arguments on (V, E).

One approach to theorems like this is to use ultraproducts: suppose that for every n, (V_n, E_n) is a graph with $\mu^2(E_n) \geq \epsilon > 0$ and $|V_n| \to \infty$.

Then in the ultraproduct $(V, E) = \prod_{\mathfrak{U}}(V_n, E_n)$, the measures $\mu^1(S) = \prod_{\mathfrak{U}} \frac{|S_n|}{|V_n|}$ and $\mu^2(T) = \prod_{\mathfrak{U}} \frac{|T_n|}{|V_n|^2}$ extend to actual σ-additive measures and $\mu^2(E) \geq \epsilon > 0$.

We can prove theorems about (V_n, E_n) for large n by using measure theoretic arguments on (V, E).

Raw measure theory usually isn't quite enough. For instance, $E \subseteq V^2$ is generally *not* measurable in the product measure. Some additional mechanism (like Keisler's notion of a *graded measure space*) is needed.

When (V_n, Γ_n) are a sequence of graphs with $|\Gamma_n| \approx \gamma |V_n|^c$, $c < 2$, we can always form an ultraproduct (V, Γ).

Question

What special properties does the ultraproduct have if the Γ_n were pseudorandom?

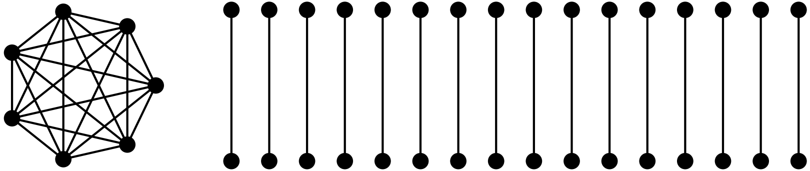

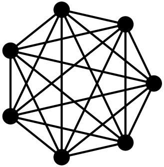

Few vertices,
Many edges

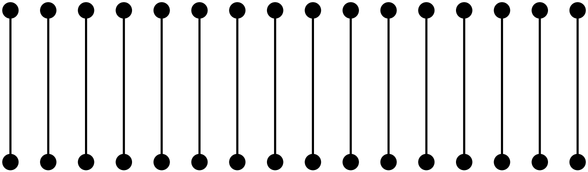

Many vertices, few edges

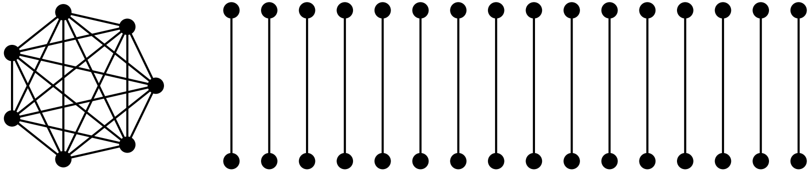

If $f(x)$ is a function on the vertices, $\int f(x)d\mu^1(x)$ is the average value of $f(x)$ over all the vertices.

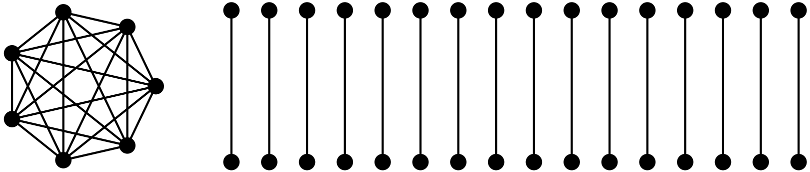

If $f(x)$ is a function on the vertices, $\int f(x) d\mu^1(x)$ is the average value of $f(x)$ over all the vertices.

If $g(x, y)$ is a function on pairs of vertices—that is, on edges—then $\int g(x, y) d\mu^2_\Gamma(x, y)$ is the average value over edges present in the graph above.

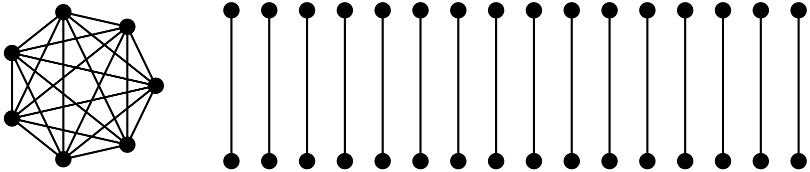

If $f(x)$ is a function on the vertices, $\int f(x) d\mu^1(x)$ is the average value of $f(x)$ over all the vertices.

If $g(x, y)$ is a function on pairs of vertices—that is, on edges—then $\int g(x, y) d\mu^2_\Gamma(x, y)$ is the average value over edges present in the graph above.

Fubini's theorem fails very badly: $\int \chi_\Gamma(x, y) d\mu^2_\Gamma(x, y) = 1$ while $\int\int \chi_\Gamma(x, y) d\mu^1(y) d\mu^1(x) \to 0$.

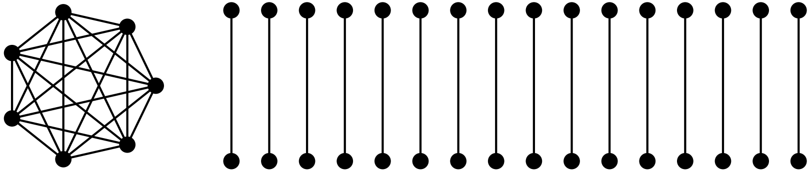

We can define additional measures: for each vertex x, define $\mu^1_x(S) = \frac{|S \cap N_\Gamma(x)|}{|N_\Gamma(x)|}$ where $N_\Gamma(x)$ is the *neighbors* of x in Γ.

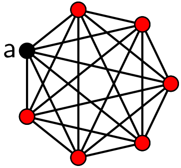

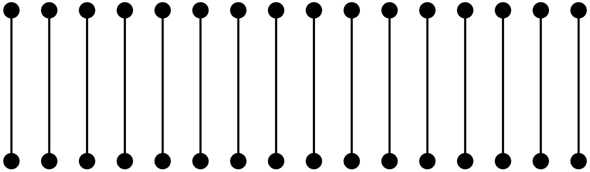

We can define additional measures: for each vertex x, define $\mu^1_x(S) = \frac{|S \cap N_\Gamma(x)|}{|N_\Gamma(x)|}$ where $N_\Gamma(x)$ is the *neighbors* of x in Γ.

In this case $\mu^1_a(S)$ concentrates on the red vertices.

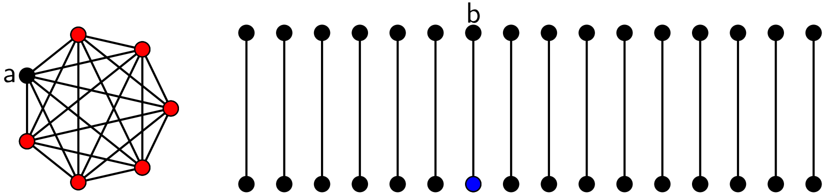

We can define additional measures: for each vertex x, define $\mu^1_x(S) = \frac{|S \cap N_\Gamma(x)|}{|N_\Gamma(x)|}$ where $N_\Gamma(x)$ is the *neighbors* of x in Γ.

In this case $\mu^1_a(S)$ concentrates on the red vertices.

$\mu^1_b(S)$ is always 0 or 1, depending on whether the blue vertex belongs to S.

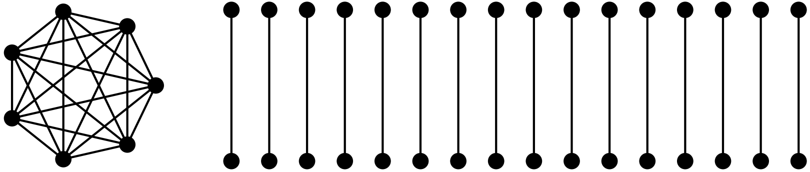

A reasonable question is the relationship between $\int g(x,y)d\mu_\Gamma^2(x,y)$ and $\int\int g(x,y)d\mu_x^1(y)d\mu^1(x)$.

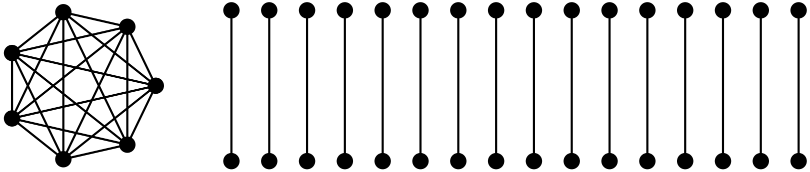

A reasonable question is the relationship between $\int g(x,y)d\mu^2_\Gamma(x,y)$ and $\int\int g(x,y)d\mu^1_x(y)d\mu^1(x)$.

$$\int \chi_\Gamma(x,y)d\mu^2_\Gamma(x,y) = 1 = \int\int \chi_\Gamma(x,y)d\mu^1_x(y)d\mu^1(x)$$

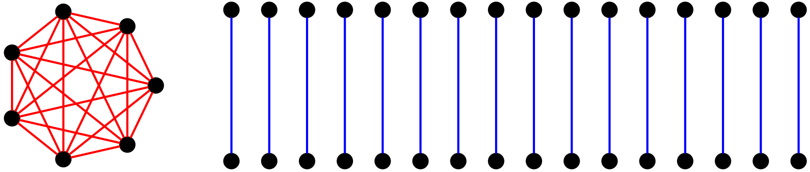

A reasonable question is the relationship between $\int g(x,y)d\mu^2_\Gamma(x,y)$ and $\int\int g(x,y)d\mu^1_x(y)d\mu^1(x)$.

Let R be the set of red edges. Then:

- $\int \chi_R(x,y)d\mu^2_\Gamma(x,y) \to 1$,
- $\int\int \chi_R(x,y)d\mu^1_x(y)d\mu^1(x) \to 0$.

A natural notion of pseudorandomness is to ask that these integrals should agree:

$$\int g(x,y)d\mu_\Gamma^2(x,y) = \iint g(x,y)d\mu_x^1(y)d\mu^1(x).$$

In finite models we only expect approximate equality up to an error term, but in the ultraproduct we expect exact equality.

A natural notion of pseudorandomness is to ask that these integrals should agree:

$$\int g(x,y)d\mu_\Gamma^2(x,y) = \iint g(x,y)d\mu_x^1(y)d\mu^1(x).$$

In finite models we only expect approximate equality up to an error term, but in the ultraproduct we expect exact equality.

These two integrals correspond to methods of selecting an edge in Γ: μ_Γ^2 selects an edge in one step while $\mu^1 \times \mu_x^1$ first selects a vertex and then extends the vertex to an edge in Γ.

More generally, let (G, E) be any fixed finite (i.e. small) graph. If we want to select a copy of (G, E) in Γ, there are many ways of doing so:

- We can select a copy of (G, E) in one step,

More generally, let (G, E) be any fixed finite (i.e. small) graph. If we want to select a copy of (G, E) in Γ, there are many ways of doing so:

- We can select a copy of (G, E) in one step,
- We can take *any* subset $H \subsetneq G$, select a copy of $(H, E \restriction [H]^2)$, and then extend this copy to a copy of (G, E),

More generally, let (G, E) be any fixed finite (i.e. small) graph. If we want to select a copy of (G, E) in Γ, there are many ways of doing so:

- We can select a copy of (G, E) in one step,
- We can take *any* subset $H \subsetneq G$, select a copy of $(H, E \upharpoonright [H]^2)$, and then extend this copy to a copy of (G, E),
- We can take any subsets $H_0 \subsetneq H_1 \subsetneq G$, select a copy of $(H_0, E \upharpoonright [H_0]^2)$, then extend this to a copy of $(H_1, E \upharpoonright [H_1]^2)$, and then extend this copy to a copy of (G, E).

More generally, let (G, E) be any fixed finite (i.e. small) graph. If we want to select a copy of (G, E) in Γ, there are many ways of doing so:

- We can select a copy of (G, E) in one step,
- We can take *any* subset $H \subsetneq G$, select a copy of $(H, E \upharpoonright [H]^2)$, and then extend this copy to a copy of (G, E),
- We can take any subsets $H_0 \subsetneq H_1 \subsetneq G$, select a copy of $(H_0, E \upharpoonright [H_0]^2)$, then extend this to a copy of $(H_1, E \upharpoonright [H_1]^2)$, and then extend this copy to a copy of (G, E).

A useful (but very strong) notion of pseudorandomness is to require that all such methods of counting give the same value. A more reasonable notion is to require this for some fixed finite list of graphs (G, E).