VC dimensions and regularity

Yuval Wigderson (ETH Zürich)

Joint with Lior Gishboliner and Asaf Shapira

BIMSA research seminar in discrete mathematics

November 4, 2025

which yonder starry sphere
Of planets and of fixed in all her wheels
Resembles nearest, mazes intricate,
Eccentric, intervolved, yet regular
Then most, when most irregular they seem;

John Milton, *Paradise Lost* V.620–4

Talk overview

Goal: Understand the regularity lemma.

Talk overview

Goal: Understand the regularity lemma.

Regularity lemma

A discrete object can be partitioned into a small number of random-like pieces.

Talk overview

Goal: Understand the regularity lemma.

Regularity lemma

A discrete object can be partitioned into a small number of random-like pieces.

Question 1

How small can small be?

Talk overview

Goal: Understand the regularity lemma.

Regularity lemma

A discrete object can be partitioned into a small number of random-like pieces.

Question 1

How small can small be?

Question 2

Can we make small smaller if we assume that the object is simple?

Talk overview

Goal: Understand the regularity lemma.

Regularity lemma

A discrete object can be partitioned into a small number of random-like pieces.

Question 1

How small can small be?

Question 2

Can we make small smaller if we assume that the object is simple?

Question 3

What does it mean for a (hyper)graph to be simple?

Regular pairs in graphs

Regular pairs in graphs

The density of a bipartite graph (X, Y) is $\frac{e(X,Y)}{|X||Y|}$.

Regular pairs in graphs

The density of a bipartite graph (X, Y) is $\frac{e(X,Y)}{|X||Y|}$.

Definition

A bipartite graph (X, Y) is ε-regular if all large induced subgraphs have basically the same density:

Regular pairs in graphs

The density of a bipartite graph (X, Y) is $\frac{e(X,Y)}{|X||Y|}$.

Definition

A bipartite graph (X, Y) is ε-regular if all large induced subgraphs have basically the same density:
For every $X' \subseteq X, Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|$, it holds that

$$|d(X', Y') - d(X, Y)| \leqslant \varepsilon.$$

Regular pairs in graphs

The density of a bipartite graph (X, Y) is $\frac{e(X,Y)}{|X||Y|}$.

Definition

A bipartite graph (X, Y) is ε-regular if all large induced subgraphs have basically the same density:
For every $X' \subseteq X, Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|$, it holds that

$$|d(X', Y') - d(X, Y)| \leqslant \varepsilon.$$

If $\varepsilon \ll d(X, Y)$, this says that (X, Y) is random-like.

Regular pairs in graphs

The density of a bipartite graph (X, Y) is $\frac{e(X,Y)}{|X||Y|}$.

Definition

A bipartite graph (X, Y) is ε-regular if all large induced subgraphs have basically the same density:
For every $X' \subseteq X, Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|$, it holds that

$$|d(X', Y') - d(X, Y)| \leqslant \varepsilon.$$

If $\varepsilon \ll d(X, Y)$, this says that (X, Y) is random-like.

If $\varepsilon \geqslant d(X, Y)$, it says basically nothing.

Regular pairs in graphs

The density of a bipartite graph (X, Y) is $\frac{e(X,Y)}{|X||Y|}$.

Definition

A bipartite graph (X, Y) is ε-regular if all large induced subgraphs have basically the same density:
For every $X' \subseteq X, Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|$, it holds that

$$|d(X', Y') - d(X, Y)| \leqslant \varepsilon.$$

If $\varepsilon \ll d(X, Y)$, this says that (X, Y) is random-like.

If $\varepsilon \geqslant d(X, Y)$, it says basically nothing.

Theorem (Induced counting lemma)

If (X, Y) is ε-regular, it contains a bi-induced copy of any fixed bipartite graph

Bi-induced: edges across (X, Y) are induced, edges inside 🤷.

Regular pairs in graphs

The density of a bipartite graph (X, Y) is $\frac{e(X,Y)}{|X||Y|}$.

Definition

A bipartite graph (X, Y) is ε-regular if all large induced subgraphs have basically the same density:
For every $X' \subseteq X, Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|$, it holds that

$$|d(X', Y') - d(X, Y)| \leqslant \varepsilon.$$

If $\varepsilon \ll d(X, Y)$, this says that (X, Y) is random-like.

If $\varepsilon \geqslant d(X, Y)$, it says basically nothing.

Theorem (Induced counting lemma)

If (X, Y) is ε-regular, it contains a bi-induced copy of any fixed bipartite graph… assuming $d(X, Y)$ is bounded away from 0 and 1.

Bi-induced: edges across (X, Y) are induced, edges inside 🤷.

Szemerédi's regularity lemma

A bipartite graph (X, Y) is ε-regular if for every $X' \subseteq X, Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|$, it holds that $|d(X', Y') - d(X, Y)| \leqslant \varepsilon$.

Szemerédi's regularity lemma

A bipartite graph (X, Y) is ε-regular if for every $X' \subseteq X, Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|$, it holds that $|d(X', Y') - d(X, Y)| \leqslant \varepsilon$.

A partition $V(G) = V_1 \sqcup \cdots \sqcup V_m$ is ε-regular if it is equitable and all but εm^2 pairs (V_i, V_j) are ε-regular.

Szemerédi's regularity lemma

A bipartite graph (X, Y) is ε-regular if for every $X' \subseteq X, Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|$, it holds that $|d(X', Y') - d(X, Y)| \leqslant \varepsilon$.

A partition $V(G) = V_1 \sqcup \cdots \sqcup V_m$ is ε-regular if it is equitable and all but εm^2 pairs (V_i, V_j) are ε-regular.

Theorem (Szemerédi)

Every graph has an ε-regular partition with $m \leqslant M(\varepsilon)$ parts.

Szemerédi's regularity lemma

A bipartite graph (X, Y) is ε-regular if for every $X' \subseteq X, Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|$, it holds that $|d(X', Y') - d(X, Y)| \leqslant \varepsilon$.

A partition $V(G) = V_1 \sqcup \cdots \sqcup V_m$ is ε-regular if it is equitable and all but εm^2 pairs (V_i, V_j) are ε-regular.

Theorem (Szemerédi)

Every graph has an ε-regular partition with $m \leqslant M(\varepsilon)$ parts.

Szemerédi: We may take $M(\varepsilon) \leqslant \text{twr}(\varepsilon^{-5})$

Szemerédi's regularity lemma

A bipartite graph (X, Y) is ε-regular if for every $X' \subseteq X, Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|$, it holds that $|d(X', Y') - d(X, Y)| \leqslant \varepsilon$.

A partition $V(G) = V_1 \sqcup \cdots \sqcup V_m$ is ε-regular if it is equitable and all but εm^2 pairs (V_i, V_j) are ε-regular.

Theorem (Szemerédi)

Every graph has an ε-regular partition with $m \leqslant M(\varepsilon)$ parts.

Szemerédi: We may take $M(\varepsilon) \leqslant \operatorname{twr}(\varepsilon^{-5}) := 2^{2^{\cdot^{\cdot^{\cdot^{2}}}}} \Big\}\, \varepsilon^{-5}$.

Szemerédi's regularity lemma

A bipartite graph (X, Y) is ε-regular if for every $X' \subseteq X, Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|$, it holds that $|d(X', Y') - d(X, Y)| \leqslant \varepsilon$.

A partition $V(G) = V_1 \sqcup \cdots \sqcup V_m$ is ε-regular if it is equitable and all but εm^2 pairs (V_i, V_j) are ε-regular.

Theorem (Szemerédi)

Every graph has an ε-regular partition with $m \leqslant M(\varepsilon)$ parts.

Szemerédi: We may take $M(\varepsilon) \leqslant \mathrm{twr}(\varepsilon^{-5}) := 2^{2^{\cdot^{\cdot^{\cdot^{2}}}}}\Big\}\, \varepsilon^{-5}$.

Gowers: Some graphs require $m \geqslant \mathrm{twr}(\varepsilon^{-1/16})$.

Szemerédi's regularity lemma

A bipartite graph (X, Y) is ε-regular if for every $X' \subseteq X, Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|$, it holds that $|d(X', Y') - d(X, Y)| \leqslant \varepsilon$.

A partition $V(G) = V_1 \sqcup \cdots \sqcup V_m$ is ε-regular if it is equitable and all but εm^2 pairs (V_i, V_j) are ε-regular.

Theorem (Szemerédi)

Every graph has an ε-regular partition with $m \leqslant M(\varepsilon)$ parts.

Szemerédi: We may take $M(\varepsilon) \leqslant \operatorname{twr}(\varepsilon^{-5}) := \left.2^{2^{\cdot^{\cdot^{\cdot^{2}}}}}\right\}\varepsilon^{-5}$.

Gowers: Some graphs require $m \geqslant \operatorname{twr}(\varepsilon^{-1/16})$.

Major question

Can we get a better bound if G is simple?

Szemerédi's regularity lemma

A bipartite graph (X, Y) is ε-regular if for every $X' \subseteq X, Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|$, it holds that $|d(X', Y') - d(X, Y)| \leqslant \varepsilon$.

A partition $V(G) = V_1 \sqcup \cdots \sqcup V_m$ is ε-regular if it is equitable and all but εm^2 pairs (V_i, V_j) are ε-regular.

Theorem (Szemerédi)

Every graph has an ε-regular partition with $m \leqslant M(\varepsilon)$ parts.

Szemerédi: We may take $M(\varepsilon) \leqslant \operatorname{twr}(\varepsilon^{-5}) := \left. 2^{2^{\cdot^{\cdot^{\cdot^{2}}}}} \right\} \varepsilon^{-5}$.

Gowers: Some graphs require $m \geqslant \operatorname{twr}(\varepsilon^{-1/16})$.

Major question

Can we get a better bound if G is simple?

Triangle-free?

Szemerédi's regularity lemma

A bipartite graph (X, Y) is ε-regular if for every $X' \subseteq X, Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|$, it holds that $|d(X', Y') - d(X, Y)| \leqslant \varepsilon$.

A partition $V(G) = V_1 \sqcup \cdots \sqcup V_m$ is ε-regular if it is equitable and all but εm^2 pairs (V_i, V_j) are ε-regular.

Theorem (Szemerédi)

Every graph has an ε-regular partition with $m \leqslant M(\varepsilon)$ parts.

Szemerédi: We may take $M(\varepsilon) \leqslant \mathrm{twr}(\varepsilon^{-5}) := 2^{2^{\cdot^{\cdot^{\cdot^{2}}}}} \Big\} \varepsilon^{-5}$.

Gowers: Some graphs require $m \geqslant \mathrm{twr}(\varepsilon^{-1/16})$.

Major question

Can we get a better bound if G is simple?

Triangle-free? No.

Szemerédi's regularity lemma

A bipartite graph (X, Y) is ε-regular if for every $X' \subseteq X, Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|$, it holds that $|d(X', Y') - d(X, Y)| \leqslant \varepsilon$.

A partition $V(G) = V_1 \sqcup \cdots \sqcup V_m$ is ε-regular if it is equitable and all but εm^2 pairs (V_i, V_j) are ε-regular.

Theorem (Szemerédi)

Every graph has an ε-regular partition with $m \leqslant M(\varepsilon)$ parts.

Szemerédi: We may take $M(\varepsilon) \leqslant \mathrm{twr}(\varepsilon^{-5}) := 2^{2^{\cdot^{\cdot^{\cdot^{2}}}}}\Big\}\,\varepsilon^{-5}$.

Gowers: Some graphs require $m \geqslant \mathrm{twr}(\varepsilon^{-1/16})$.

Major question

Can we get a better bound if G is simple?

Triangle-free? No. Bipartite?

Szemerédi's regularity lemma

A bipartite graph (X, Y) is ε-regular if for every $X' \subseteq X, Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|$, it holds that $|d(X', Y') - d(X, Y)| \leqslant \varepsilon$.

A partition $V(G) = V_1 \sqcup \cdots \sqcup V_m$ is ε-regular if it is equitable and all but εm^2 pairs (V_i, V_j) are ε-regular.

Theorem (Szemerédi)

Every graph has an ε-regular partition with $m \leqslant M(\varepsilon)$ parts.

Szemerédi: We may take $M(\varepsilon) \leqslant \mathrm{twr}(\varepsilon^{-5}) := \left.2^{2^{\cdot^{\cdot^{\cdot^{2}}}}}\right\}\varepsilon^{-5}$.

Gowers: Some graphs require $m \geqslant \mathrm{twr}(\varepsilon^{-1/16})$.

Major question

Can we get a better bound if G is simple?

Triangle-free? No. Bipartite? No.

Szemerédi's regularity lemma

A bipartite graph (X, Y) is ε-regular if for every $X' \subseteq X, Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|$, it holds that $|d(X', Y') - d(X, Y)| \leqslant \varepsilon$.

A partition $V(G) = V_1 \sqcup \cdots \sqcup V_m$ is ε-regular if it is equitable and all but εm^2 pairs (V_i, V_j) are ε-regular.

Theorem (Szemerédi)

Every graph has an ε-regular partition with $m \leqslant M(\varepsilon)$ parts.

Szemerédi: We may take $M(\varepsilon) \leqslant \operatorname{twr}(\varepsilon^{-5}) := \left.2^{2^{\cdot^{\cdot^{\cdot^{2}}}}}\right\}\varepsilon^{-5}$.

Gowers: Some graphs require $m \geqslant \operatorname{twr}(\varepsilon^{-1/16})$.

Major question

Can we get a better bound if G is simple?

Triangle-free? No. Bipartite? No. Chordal?

Szemerédi's regularity lemma

A bipartite graph (X, Y) is ε-regular if for every $X' \subseteq X$, $Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|$, $|Y'| \geqslant \varepsilon|Y|$, it holds that $|d(X', Y') - d(X, Y)| \leqslant \varepsilon$.

A partition $V(G) = V_1 \sqcup \cdots \sqcup V_m$ is ε-regular if it is equitable and all but εm^2 pairs (V_i, V_j) are ε-regular.

Theorem (Szemerédi)

Every graph has an ε-regular partition with $m \leqslant M(\varepsilon)$ parts.

Szemerédi: We may take $M(\varepsilon) \leqslant \operatorname{twr}(\varepsilon^{-5}) \coloneqq \left.2^{2^{\cdot^{\cdot^{\cdot^{2}}}}}\right\}\varepsilon^{-5}$.

Gowers: Some graphs require $m \geqslant \operatorname{twr}(\varepsilon^{-1/16})$.

Major question

Can we get a better bound if G is simple?

Triangle-free? No. Bipartite? No. Chordal? No.

Szemerédi's regularity lemma

A bipartite graph (X, Y) is ε-regular if for every $X' \subseteq X, Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|$, it holds that $|d(X', Y') - d(X, Y)| \leqslant \varepsilon$.

A partition $V(G) = V_1 \sqcup \cdots \sqcup V_m$ is ε-regular if it is equitable and all but εm^2 pairs (V_i, V_j) are ε-regular.

Theorem (Szemerédi)

Every graph has an ε-regular partition with $m \leqslant M(\varepsilon)$ parts.

Szemerédi: We may take $M(\varepsilon) \leqslant \mathrm{twr}(\varepsilon^{-5}) := 2^{2^{\cdot^{\cdot^{\cdot^{2}}}}} \Big\} \varepsilon^{-5}$.

Gowers: Some graphs require $m \geqslant \mathrm{twr}(\varepsilon^{-1/16})$.

Major question

Can we get a better bound if G is simple?

Triangle-free? No. Bipartite? No. Chordal? No. Perfect?

Szemerédi's regularity lemma

A bipartite graph (X, Y) is ε-regular if for every $X' \subseteq X, Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|$, it holds that $|d(X', Y') - d(X, Y)| \leqslant \varepsilon$.

A partition $V(G) = V_1 \sqcup \cdots \sqcup V_m$ is ε-regular if it is equitable and all but εm^2 pairs (V_i, V_j) are ε-regular.

Theorem (Szemerédi)

Every graph has an ε-regular partition with $m \leqslant M(\varepsilon)$ parts.

Szemerédi: We may take $M(\varepsilon) \leqslant \mathrm{twr}(\varepsilon^{-5}) := 2^{2^{\cdot^{\cdot^{\cdot^{2}}}}} \Big\}\, \varepsilon^{-5}$.

Gowers: Some graphs require $m \geqslant \mathrm{twr}(\varepsilon^{-1/16})$.

Major question

Can we get a better bound if G is simple?

Triangle-free? No. Bipartite? No. Chordal? No. Perfect? No.

Szemerédi's regularity lemma

A bipartite graph (X, Y) is ε-regular if for every $X' \subseteq X, Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|$, it holds that $|d(X', Y') - d(X, Y)| \leqslant \varepsilon$.

A partition $V(G) = V_1 \sqcup \cdots \sqcup V_m$ is ε-regular if it is equitable and all but εm^2 pairs (V_i, V_j) are ε-regular.

Theorem (Szemerédi)

Every graph has an ε-regular partition with $m \leqslant M(\varepsilon)$ parts.

Szemerédi: We may take $M(\varepsilon) \leqslant \text{twr}(\varepsilon^{-5}) := 2^{2^{\cdot^{\cdot^{\cdot^{2}}}}} \Big\} \varepsilon^{-5}$.

Gowers: Some graphs require $m \geqslant \text{twr}(\varepsilon^{-1/16})$.

Major question

Can we get a better bound if G is simple?

Triangle-free? No. Bipartite? No. Chordal? No. Perfect? No.

Geometrically defined?

Szemerédi's regularity lemma

A bipartite graph (X, Y) is ε-regular if for every $X' \subseteq X$, $Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|$, $|Y'| \geqslant \varepsilon|Y|$, it holds that $|d(X', Y') - d(X, Y)| \leqslant \varepsilon$.

A partition $V(G) = V_1 \sqcup \cdots \sqcup V_m$ is ε-regular if it is equitable and all but εm^2 pairs (V_i, V_j) are ε-regular.

Theorem (Szemerédi)

Every graph has an ε-regular partition with $m \leqslant M(\varepsilon)$ parts.

Szemerédi: We may take $M(\varepsilon) \leqslant \operatorname{twr}(\varepsilon^{-5}) := \left.2^{2^{\cdot^{\cdot^{\cdot^{2}}}}}\right\}\varepsilon^{-5}$.

Gowers: Some graphs require $m \geqslant \operatorname{twr}(\varepsilon^{-1/16})$.

Major question

Can we get a better bound if G is simple?

Triangle-free? No. Bipartite? No. Chordal? No. Perfect? No.

Geometrically defined? **Yes!**

VC dimension

Geometrically defined graphs have bounded Vapnik–Chervonenkis (VC) dimension.

VC dimension

Geometrically defined graphs have bounded Vapnik–Chervonenkis (VC) dimension.

Non-standard definition

If a bipartite graph B is not a bi-induced subgraph of G, then G has VC dimension $\leqslant |B|$.

VC dimension

Geometrically defined graphs have bounded Vapnik–Chervonenkis (VC) dimension.

Non-standard definition

If a bipartite graph B is not a bi-induced subgraph of G, then G has VC dimension $\leqslant |B|$.
That is, G has bounded VC dimension if it forbids some fixed bi-induced B.

VC dimension

Geometrically defined graphs have bounded Vapnik–Chervonenkis (VC) dimension.

Non-standard definition

If a bipartite graph B is not a bi-induced subgraph of G, then G has VC dimension $\leqslant |B|$.
That is, G has bounded VC dimension if it forbids some fixed bi-induced B.

VC dimension is a fundamental combinatorial notion capturing "low complexity" of discrete objects.

VC dimension

Geometrically defined graphs have bounded Vapnik–Chervonenkis (VC) dimension.

Non-standard definition

If a bipartite graph B is not a bi-induced subgraph of G, then G has VC dimension $\leqslant |B|$.
That is, G has bounded VC dimension if it forbids some fixed bi-induced B.

VC dimension is a fundamental combinatorial notion capturing “low complexity” of discrete objects. It is of central importance in learning theory, model theory, group theory, combinatorics…

VC dimension

Geometrically defined graphs have bounded Vapnik–Chervonenkis (VC) dimension.

Non-standard definition

If a bipartite graph B is not a bi-induced subgraph of G, then G has VC dimension $\leqslant |B|$.
That is, G has bounded VC dimension if it forbids some fixed bi-induced B.

VC dimension is a fundamental combinatorial notion capturing "low complexity" of discrete objects. It is of central importance in learning theory, model theory, group theory, combinatorics…

Many combinatorial questions become much simpler when restricted to graphs of bounded VC dimension.

Graph regularity and VC dimension

G has bounded VC dimension if it forbids some fixed bi-induced B.

Graph regularity and VC dimension

G has bounded VC dimension if it forbids some fixed bi-induced B.

If (X, Y) is ε-regular, it contains a bi-induced copy of any fixed bipartite graph, assuming $d(X, Y)$ is bounded away from 0 and 1.

Graph regularity and VC dimension

G has bounded VC dimension if it forbids some fixed bi-induced B.

If (X, Y) is ε-regular, it contains a bi-induced copy of any fixed bipartite graph, assuming $d(X, Y)$ is bounded away from 0 and 1.

Corollary

If G has bounded VC dimension, then every ε-regular pair in G must have density close to 0 *or* 1.

Graph regularity and VC dimension

G has bounded VC dimension if it forbids some fixed bi-induced B.

If (X, Y) is ε-regular, it contains a bi-induced copy of any fixed bipartite graph, assuming $d(X, Y)$ is bounded away from 0 and 1.

Corollary

If G has bounded VC dimension, then every ε-regular pair in G must have density close to 0 *or* 1.

Definition

A pair (X, Y) is ε-homogeneous if $d(X, Y) \in [0, \varepsilon] \cup [1 - \varepsilon, 1]$.

Graph regularity and VC dimension

G has bounded VC dimension if it forbids some fixed bi-induced B.

If (X, Y) is ε-regular, it contains a bi-induced copy of any fixed bipartite graph, assuming $d(X, Y)$ is bounded away from 0 and 1.

Corollary

If G has bounded VC dimension, then every ε-regular pair in G must have density close to 0 or 1.

Definition

A pair (X, Y) is ε-homogeneous if $d(X, Y) \in [0, \varepsilon] \cup [1 - \varepsilon, 1]$.
A partition $V(G) = V_1 \sqcup \cdots \sqcup V_m$ is ε-homogeneous if it is equitable and all but εm^2 pairs are ε-homogeneous.

Graph regularity and VC dimension

G has bounded VC dimension if it forbids some fixed bi-induced B.

If (X, Y) is ε-regular, it contains a bi-induced copy of any fixed bipartite graph, assuming $d(X, Y)$ is bounded away from 0 and 1.

Corollary

If G has bounded VC dimension, then every ε-regular pair in G must have density close to 0 or 1.

Definition

A pair (X, Y) is ε-homogeneous if $d(X, Y) \in [0, \varepsilon] \cup [1 - \varepsilon, 1]$.
A partition $V(G) = V_1 \sqcup \cdots \sqcup V_m$ is ε-homogeneous if it is equitable and all but εm^2 pairs are ε-homogeneous.

Corollary (Regularity lemma for bounded VC dimension)

If G has bounded VC dimension, it has an ε-homogeneous partition with $m \leqslant M(\varepsilon)$ parts.

Quantitative bounds

Szemerédi, Gowers: Every graph G has an ε-regular partition into $m \leqslant \mathrm{twr}(\varepsilon^{-5})$ parts, and for some G such a bound is necessary.

Quantitative bounds

Szemerédi, Gowers: Every graph G has an ε-regular partition into $m \leqslant \mathrm{twr}(\varepsilon^{-5})$ parts, and for some G such a bound is necessary.

Theorem (Alon–Fischer–Newman, Lovász–Szegedy)

If G has bounded VC dimension, it has an ε-homogeneous partition into $m \leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts (and this is best possible).

Quantitative bounds

Szemerédi, Gowers: Every graph G has an ε-regular partition into $m \leqslant \mathrm{twr}(\varepsilon^{-5})$ parts, and for some G such a bound is necessary.

Theorem (Alon–Fischer–Newman, Lovász–Szegedy)

If G has bounded VC dimension, it has an ε-homogeneous partition into $m \leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts (and this is best possible).

Theorem

The following are equivalent for a hereditary family of graphs $\mathcal{G}$.

Quantitative bounds

Szemerédi, Gowers: Every graph G has an ε-regular partition into $m \leqslant \mathrm{twr}(\varepsilon^{-5})$ parts, and for some G such a bound is necessary.

Theorem (Alon–Fischer–Newman, Lovász–Szegedy)

If G has bounded VC dimension, it has an ε-homogeneous partition into $m \leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts (and this is best possible).

Theorem

The following are equivalent for a hereditary family of graphs $\mathcal{G}$.

- *$\mathcal{G}$ has bounded VC dimension.*

Quantitative bounds

Szemerédi, Gowers: Every graph G has an ε-regular partition into $m \leqslant \operatorname{twr}(\varepsilon^{-5})$ parts, and for some G such a bound is necessary.

Theorem (Alon–Fischer–Newman, Lovász–Szegedy)

If G has bounded VC dimension, it has an ε-homogeneous partition into $m \leqslant \operatorname{poly}(\frac{1}{\varepsilon})$ parts (and this is best possible).

Theorem

The following are equivalent for a hereditary family of graphs $\mathcal{G}$.

- *$\mathcal{G}$ has bounded VC dimension.*
- *$\mathcal{G}$ has an ε-homogeneous partition with $\leqslant \operatorname{poly}(\frac{1}{\varepsilon})$ parts.*

Quantitative bounds

Szemerédi, Gowers: Every graph G has an ε-regular partition into $m \leqslant \text{twr}(\varepsilon^{-5})$ parts, and for some G such a bound is necessary.

Theorem (Alon–Fischer–Newman, Lovász–Szegedy)

If G has bounded VC dimension, it has an ε-homogeneous partition into $m \leqslant \text{poly}(\frac{1}{\varepsilon})$ parts (and this is best possible).

Theorem

The following are equivalent for a hereditary family of graphs $\mathcal{G}$.

- *$\mathcal{G}$ has bounded VC dimension.*
- *$\mathcal{G}$ has an ε-homogeneous partition with $\leqslant \text{poly}(\frac{1}{\varepsilon})$ parts.*
- *$\mathcal{G}$ has an ε-homogeneous partition with any number of parts.*

Quantitative bounds

Szemerédi, Gowers: Every graph G has an ε-regular partition into $m \leqslant \mathrm{twr}(\varepsilon^{-5})$ parts, and for some G such a bound is necessary.

Theorem (Alon–Fischer–Newman, Lovász–Szegedy)

If G has bounded VC dimension, it has an ε-homogeneous partition into $m \leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts (and this is best possible).

Theorem

The following are equivalent for a hereditary family of graphs $\mathcal{G}$.

- *$\mathcal{G}$ has bounded VC dimension.*
- *$\mathcal{G}$ has an ε-homogeneous partition with $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts.*
- *$\mathcal{G}$ has an ε-homogeneous partition with any number of parts.*
- *$\mathcal{G}$ has an ε-regular partition with $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts.*

Quantitative bounds

Szemerédi, Gowers: Every graph G has an ε-regular partition into $m \leqslant \operatorname{twr}(\varepsilon^{-5})$ parts, and for some G such a bound is necessary.

Theorem (Alon–Fischer–Newman, Lovász–Szegedy)

If G has bounded VC dimension, it has an ε-homogeneous partition into $m \leqslant \operatorname{poly}(\frac{1}{\varepsilon})$ parts (and this is best possible).

Theorem

The following are equivalent for a hereditary family of graphs $\mathcal{G}$.

- *$\mathcal{G}$ has bounded VC dimension.*
- *$\mathcal{G}$ has an ε-homogeneous partition with $\leqslant \operatorname{poly}(\frac{1}{\varepsilon})$ parts.*
- *$\mathcal{G}$ has an ε-homogeneous partition with any number of parts.*
- *$\mathcal{G}$ has an ε-regular partition with $\leqslant \operatorname{poly}(\frac{1}{\varepsilon})$ parts.*
- *$\mathcal{G}$ has an ε-regular partition with a sub-tower number of parts.*

Why are hypergraphs harder?

Theorem

The following are equivalent for a hereditary family of graphs G.

- *G has bounded VC dimension.*
- *G has an ε-homogeneous partition with* $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ *parts.*
- *G has an ε-homogeneous partition with any number of parts.*
- *G has an ε-regular partition with* $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ *parts.*
- *G has an ε-regular partition with a sub-tower number of parts.*

Why are hypergraphs harder?

Theorem

The following are equivalent for a hereditary family of graphs G.

- *G has bounded VC dimension.*
- *G has an ε-homogeneous partition with $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts.*
- *G has an ε-homogeneous partition with any number of parts.*
- *G has an ε-regular partition with $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts.*
- *G has an ε-regular partition with a sub-tower number of parts.*

We seek a hypergraph generalization of this theorem.

Why are hypergraphs harder?

Theorem

The following are equivalent for a hereditary family of graphs G.

- *G has bounded VC dimension.*
- *G has an ε-homogeneous partition with $\leqslant \text{poly}(\frac{1}{\varepsilon})$ parts.*
- *G has an ε-homogeneous partition with any number of parts.*
- *G has an ε-regular partition with $\leqslant \text{poly}(\frac{1}{\varepsilon})$ parts.*
- *G has an ε-regular partition with a sub-tower number of parts.*

We seek a hypergraph generalization of this theorem.

In k-uniform hypergraphs, there are $\geqslant 2$ notions of regularity

Why are hypergraphs harder?

Theorem

The following are equivalent for a hereditary family of graphs G.

- *G has bounded VC dimension.*
- *G has an ε-homogeneous partition with $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts.*
- *G has an ε-homogeneous partition with any number of parts.*
- *G has an ε-regular partition with $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts.*
- *G has an ε-regular partition with a sub-tower number of parts.*

We seek a hypergraph generalization of this theorem.

In k-uniform hypergraphs, there are $\geqslant 2$ notions of regularity... and $\geqslant k$ notions of VC dimension.

Why are hypergraphs harder?

Theorem

The following are equivalent for a hereditary family of graphs G.

- *G has bounded VC dimension.*
- *G has an ε-homogeneous partition with* $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ *parts.*
- *G has an ε-homogeneous partition with any number of parts.*
- *G has an ε-regular partition with* $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ *parts.*
- *G has an ε-regular partition with a sub-tower number of parts.*

We seek a hypergraph generalization of this theorem.

In k-uniform hypergraphs, there are $\geqslant 2$ notions of regularity… and $\geqslant k$ notions of VC dimension.

Also, the regularity notions can be very hard to work with.

A statement of the hypergraph regularity lemma

Theorem 11 (Hypergraph Regularity Lemma). *For all positive reals* μ *and* δ_k *and functions*

$$\delta_j\colon (0,1]^{k-j} \to (0,1] \quad \textit{for } j = 2, \ldots, k-1,$$

$$\textit{and } r\colon \mathbb{N} \times (0,1]^{k-2} \to \mathbb{N},$$

there exist T_0 *and* n_0 *so that the following holds. For every k-graph* $\mathcal{H}^{(k)}$ *on* $n \geq n_0$ *vertices, there exist a family of partitions* $\mathcal{P} = \mathcal{P}(k-1, \mathbf{a})$ *and a vector* $\mathbf{d} = (d_2, \ldots, d_{k-1})$ *so that, for* $\boldsymbol{\delta} = (\delta_2, \ldots, \delta_{k-1})$*, where* $\delta_j = \delta_j(d_j, \ldots, d_{k-1})$ *for all j, and* $r = r(a_1, \mathbf{d})$*, the following holds:*

(*i*) $\mathcal{P}$ *is a* $(\mu, \boldsymbol{\delta}, \mathbf{d}, r)$*-equitable family of partitions and* $a_i \leq T_0$ *for every* $i = 1, \ldots, k-1$ *and*
(*ii*) $\mathcal{H}^{(k)}$ *is* (δ_k, r)*-regular w.r.t.* $\mathcal{P}$.

[Rödl–Nagle–Skokan–Schacht–Kohayakawa]

A statement of the hypergraph regularity lemma

Theorem 11 (Hypergraph Regularity Lemma). *For all positive reals* μ *and* δ_k *and functions*

$$\delta_j\colon (0,1]^{k-j} \to (0,1] \quad \text{for } j = 2, \ldots, k-1,$$

$$\text{and } r\colon \mathbb{N} \times (0,1]^{k-2} \to \mathbb{N},$$

there exist T_0 *and* n_0 *so that the following holds. For every k-graph* $\mathcal{H}^{(k)}$ *on* $n \geq n_0$ *vertices, there exist a family of partitions* $\mathcal{P} = \mathcal{P}(k-1, \mathbf{a})$ *and a vector* $\mathbf{d} = (d_2, \ldots, d_{k-1})$ *so that, for* $\boldsymbol{\delta} = (\delta_2, \ldots, \delta_{k-1})$*, where* $\delta_j = \delta_j(d_j, \ldots, d_{k-1})$ *for all j, and* $r = r(a_1, \mathbf{d})$*, the following holds:*

(i) $\mathcal{P}$ *is a* $(\mu, \boldsymbol{\delta}, \mathbf{d}, r)$*-equitable family of partitions and* $a_i \leq T_0$ *for every* $i = 1, \ldots, k-1$ *and*
(ii) $\mathcal{H}^{(k)}$ *is* (δ_k, r)*-regular w.r.t.* $\mathcal{P}$.

[Rödl–Nagle–Skokan–Schacht–Kohayakawa]

Hypergraphs are scary! Let's try to only think about graphs.

Weak regularity for hypergraphs

Weak regularity for hypergraphs

The density of a tripartite 3-graph (X, Y, Z) is $\frac{e(X,Y,Z)}{|X||Y||Z|}$.

Weak regularity for hypergraphs

The density of a tripartite 3-graph (X, Y, Z) is $\frac{e(X,Y,Z)}{|X||Y||Z|}$.

It is weakly ε-regular if $|d(X', Y', Z') - d(X, Y, Z)| \leqslant \varepsilon$ whenever $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|, |Z'| \geqslant \varepsilon|Z|$.

Weak regularity for hypergraphs

The density of a tripartite 3-graph (X, Y, Z) is $\frac{e(X,Y,Z)}{|X||Y||Z|}$.

It is weakly ε-regular if $|d(X', Y', Z') - d(X, Y, Z)| \leqslant \varepsilon$ whenever $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|, |Z'| \geqslant \varepsilon|Z|$.

Theorem (Chung)

Every 3-graph has a weakly ε-regular partition into $m \leqslant M(\varepsilon)$ parts.

Weak regularity for hypergraphs

The density of a tripartite 3-graph (X, Y, Z) is $\frac{e(X,Y,Z)}{|X||Y||Z|}$.

It is weakly ε-regular if $|d(X', Y', Z') - d(X, Y, Z)| \leqslant \varepsilon$ whenever $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|, |Z'| \geqslant \varepsilon|Z|$.

Theorem (Chung)

Every 3-graph has a weakly ε-regular partition into $m \leqslant M(\varepsilon)$ parts.

Chung's proof gives $M(\varepsilon) \leqslant \mathrm{twr}(\varepsilon^{-C})$, and Gowers's result shows this is necessary in general.

Weak regularity for hypergraphs

The density of a tripartite 3-graph (X, Y, Z) is $\frac{e(X,Y,Z)}{|X||Y||Z|}$.

It is weakly ε-regular if $|d(X', Y', Z') - d(X, Y, Z)| \leqslant \varepsilon$ whenever $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|, |Z'| \geqslant \varepsilon|Z|$.

Theorem (Chung)

Every 3-graph has a weakly ε-regular partition into $m \leqslant M(\varepsilon)$ parts.

Chung's proof gives $M(\varepsilon) \leqslant \mathrm{twr}(\varepsilon^{-C})$, and Gowers's result shows this is necessary in general.

Bad news: This notion of regularity is too weak for most applications, because it does not support a counting lemma.

Weak regularity for hypergraphs

The density of a tripartite 3-graph (X, Y, Z) is $\frac{e(X,Y,Z)}{|X||Y||Z|}$.

It is weakly ε-regular if $|d(X', Y', Z') - d(X, Y, Z)| \leqslant \varepsilon$ whenever $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|, |Z'| \geqslant \varepsilon|Z|$.

Theorem (Chung)

Every 3-graph has a weakly ε-regular partition into $m \leqslant M(\varepsilon)$ parts.

Chung's proof gives $M(\varepsilon) \leqslant \mathrm{twr}(\varepsilon^{-C})$, and Gowers's result shows this is necessary in general.

Bad news: This notion of regularity is too weak for most applications, because it does not support a counting lemma.

Nonetheless, it's still interesting, and useful in some settings.

Weak regularity for hypergraphs

The density of a tripartite 3-graph (X, Y, Z) is $\frac{e(X,Y,Z)}{|X||Y||Z|}$.

It is weakly ε-regular if $|d(X', Y', Z') - d(X, Y, Z)| \leqslant \varepsilon$ whenever $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|, |Z'| \geqslant \varepsilon|Z|$.

Theorem (Chung)

Every 3-graph has a weakly ε-regular partition into $m \leqslant M(\varepsilon)$ parts.

Chung's proof gives $M(\varepsilon) \leqslant \mathrm{twr}(\varepsilon^{-C})$, and Gowers's result shows this is necessary in general.

Bad news: This notion of regularity is too weak for most applications, because it does not support a counting lemma.

Nonetheless, it's still interesting, and useful in some settings.

Question

What does it mean for a hypergraph to be simple? Can we strengthen Chung's regularity lemma for simple hypergraphs?

Weak regularity for hypergraphs

The density of a tripartite 3-graph (X, Y, Z) is $\frac{e(X,Y,Z)}{|X||Y||Z|}$.

It is weakly ε-regular if $|d(X', Y', Z') - d(X, Y, Z)| \leqslant \varepsilon$ whenever $|X'| \geqslant \varepsilon|X|, |Y'| \geqslant \varepsilon|Y|, |Z'| \geqslant \varepsilon|Z|$.

Theorem (Chung)

Every 3-graph has a weakly ε-regular partition into $m \leqslant M(\varepsilon)$ parts.

Chung's proof gives $M(\varepsilon) \leqslant \mathrm{twr}(\varepsilon^{-C})$, and Gowers's result shows this is necessary in general.

Bad news: This notion of regularity is too weak for most applications, because it does not support a counting lemma.

Nonetheless, it's still interesting, and useful in some settings.

Question

What does it mean for a hypergraph to be simple? Can we strengthen Chung's regularity lemma for simple hypergraphs?

We'll see three different answers to this question.

Strong VC dimension

Question

What does it mean for a 3-graph to be simple?

Strong VC dimension

Question

What does it mean for a 3-graph to be simple?

Hypergraphs are scary! Let's try to only think about graphs.

Strong VC dimension

Question

What does it mean for a 3-graph to be simple?

Hypergraphs are scary! Let's try to only think about graphs.

Geometrically defined graphs have bounded VC dimension.

Theorem (Alon–Fischer–Newman, Lovász–Szegedy)

If G has bounded VC dimension, it has an ε-homogeneous partition into $m \leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts (and this is best possible).

Strong VC dimension

Question

What does it mean for a 3-graph to be simple?

Hypergraphs are scary! Let's try to only think about graphs.

Geometrically defined graphs have bounded VC dimension.

Theorem (Alon–Fischer–Newman, Lovász–Szegedy)

If G has bounded VC dimension, it has an ε-homogeneous partition into $m \leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts (and this is best possible).

Geometrically defined hypergraphs have bounded strong VC dimension

Strong VC dimension

Question

What does it mean for a 3-graph to be simple?

Hypergraphs are scary! Let's try to only think about graphs.

Geometrically defined graphs have bounded VC dimension.

Theorem (Alon–Fischer–Newman, Lovász–Szegedy)

If G has bounded VC dimension, it has an ε-homogeneous partition into $m \leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts (and this is best possible).

Geometrically defined hypergraphs have bounded strong VC dimension... which I won't define.

Strong VC dimension

Question

What does it mean for a 3-graph to be simple?

Hypergraphs are scary! Let's try to only think about graphs.

Geometrically defined graphs have bounded VC dimension.

Theorem (Alon–Fischer–Newman, Lovász–Szegedy)

If G has bounded VC dimension, it has an ε-homogeneous partition into $m \leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts (and this is best possible).

Geometrically defined hypergraphs have bounded strong VC dimension… which I won't define. It arises naturally from geometry.

Strong VC dimension

Question

What does it mean for a 3-graph to be simple?

Hypergraphs are scary! Let's try to only think about graphs.

Geometrically defined graphs have bounded VC dimension.

Theorem (Alon–Fischer–Newman, Lovász–Szegedy)

If G has bounded VC dimension, it has an ε-homogeneous partition into $m \leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts (and this is best possible).

Geometrically defined hypergraphs have bounded strong VC dimension… which I won't define. It arises naturally from geometry.

Theorem (Chernikov–Starchenko, Fox–Pach–Suk)

If H has bounded strong VC dimension, it has an ε-homogeneous partition into $m \leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts (and this is best possible).

Is it a characterization?

Theorem (Chernikov–Starchenko, Fox–Pach–Suk)

If H has bounded strong VC dimension, it has an ε-homogeneous partition into $m \leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts (and this is best possible).

Is it a characterization?

Theorem (Chernikov–Starchenko, Fox–Pach–Suk)

If H has bounded strong VC dimension, it has an ε-homogeneous partition into $m \leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts (and this is best possible).

Theorem (Terry)

The following are equivalent for a hereditary family of 3-graphs H.

Is it a characterization?

Theorem (Chernikov–Starchenko, Fox–Pach–Suk)

If H has bounded strong VC dimension, it has an ε-homogeneous partition into $m \leqslant \text{poly}(\frac{1}{\varepsilon})$ parts (and this is best possible).

Theorem (Terry)

The following are equivalent for a hereditary family of 3-graphs H.

- *H has bounded strong VC dimension.*

Is it a characterization?

Theorem (Chernikov–Starchenko, Fox–Pach–Suk)

If H has bounded strong VC dimension, it has an ε-homogeneous partition into $m \leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts (and this is best possible).

Theorem (Terry)

The following are equivalent for a hereditary family of 3-graphs H.

- *H has bounded strong VC dimension.*
- *H has an ε-homogeneous partition with $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts.*

Is it a characterization?

Theorem (Chernikov–Starchenko, Fox–Pach–Suk)

If H has bounded strong VC dimension, it has an ε-homogeneous partition into $m \leqslant \text{poly}(\frac{1}{\varepsilon})$ parts (and this is best possible).

Theorem (Terry)

The following are equivalent for a hereditary family of 3-graphs H.

- *H has bounded strong VC dimension.*
- *H has an ε-homogeneous partition with $\leqslant \text{poly}(\frac{1}{\varepsilon})$ parts.*
- *H has a weakly ε-regular partition with $\leqslant \text{poly}(\frac{1}{\varepsilon})$ parts.*

Is it a characterization?

Theorem (Chernikov–Starchenko, Fox–Pach–Suk)

If H has bounded strong VC dimension, it has an ε-homogeneous partition into $m \leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts (and this is best possible).

Theorem (Terry)

The following are equivalent for a hereditary family of 3-graphs H.

- *H has bounded strong VC dimension.*
- *H has an ε-homogeneous partition with $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts.*
- *H has a weakly ε-regular partition with $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts.*

However, strong VC dimension does not characterize the existence of homogeneous partitions, nor of having sub-tower bounds.

Is it a characterization?

Theorem (Chernikov–Starchenko, Fox–Pach–Suk)

If H has bounded strong VC dimension, it has an ε-homogeneous partition into $m \leqslant \text{poly}(\frac{1}{\varepsilon})$ parts (and this is best possible).

Theorem (Terry)

The following are equivalent for a hereditary family of 3-graphs H.

- *H has bounded strong VC dimension.*
- *H has an ε-homogeneous partition with $\leqslant \text{poly}(\frac{1}{\varepsilon})$ parts.*
- *H has a weakly ε-regular partition with $\leqslant \text{poly}(\frac{1}{\varepsilon})$ parts.*

However, strong VC dimension does not characterize the existence of homogeneous partitions, nor of having sub-tower bounds.

Time to come up with some other notions!

Simple links

Question

What does it mean for a 3-graph to be simple?

Simple links

Question

What does it mean for a 3-graph to be simple?

Hypergraphs are scary! Let's try to only think about graphs.

Simple links

Question

What does it mean for a 3-graph to be simple?

Hypergraphs are scary! Let's try to only think about graphs.

Link of x: Graph with edge set $\{yz : xyz \in E(H)\}$.

Simple links

Question

What does it mean for a 3-graph to be simple?

Hypergraphs are scary! Let's try to only think about graphs.

Link of x: Graph with edge set $\{yz : xyz \in E(H)\}$.

"Definition"

A 3-graph H is simple if the link of every vertex is simple.

Simple links

Question

What does it mean for a 3-graph to be simple?

Hypergraphs are scary! Let's try to only think about graphs.

Link of x: Graph with edge set $\{yz : xyz \in E(H)\}$.

"Definition"

A 3-graph H is simple if the link of every vertex is simple.

This "definition" is of central importance in the theory of high-dimensional expanders, and shows up in many Ramsey- and Turán-type questions in hypergraphs.

Do simple links help with regularity?

"Definition"

A 3-graph H is simple if the link of every vertex is simple.

Do simple links help with regularity?

"Definition"

A 3-graph H is simple if the link of every vertex is simple.

Suppose I promise that every link has a small ε-regular partition.

Do simple links help with regularity?

"Definition"

A 3-graph H is simple if the link of every vertex is simple.

Suppose I promise that every link has a small ε-regular partition.
Does this help you find a weakly ε-regular partition of H?

Do simple links help with regularity?

"Definition"

A 3-graph H is simple if the link of every vertex is simple.

Suppose I promise that every link has a small ε-regular partition.
Does this help you find a weakly ε-regular partition of H?

No!

Yes!

Do simple links help with regularity?

"Definition"

A 3-graph H is simple if the link of every vertex is simple.

Suppose I promise that every link has a small ε-regular partition.
Does this help you find a weakly ε-regular partition of H?

No!

Theorem (Gishboliner–Shapira–W.)

There exists a 3-graph H such that

- *every link has an ε-regular partition with* $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ *parts, but*

Yes!

Do simple links help with regularity?

"Definition"

A 3-graph H is simple if the link of every vertex is simple.

Suppose I promise that every link has a small ε-regular partition.
Does this help you find a weakly ε-regular partition of H?

No!

Theorem (Gishboliner–Shapira–W.)

There exists a 3-graph H such that

- *every link has an ε-regular partition with $\leqslant \text{poly}(\frac{1}{\varepsilon})$ parts, but*
- *every weakly ε-regular partition of H has $\geqslant \text{twr}(\log \frac{1}{\varepsilon})$ parts.*

Yes!

Do simple links help with regularity?

"Definition"

A 3-graph H is simple if the link of every vertex is simple.

Suppose I promise that every link has a small ε-regular partition.
Does this help you find a weakly ε-regular partition of H?

No!

Theorem (Gishboliner–Shapira–W.)

There exists a 3-graph H such that

- *every link has an ε-regular partition with $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts, but*
- *every weakly ε-regular partition of H has $\geqslant \mathrm{twr}(\log \frac{1}{\varepsilon})$ parts.*

Yes!

Theorem (Gishboliner–Shapira–W.)

If every link of H has an ε-homogeneous partition with $\leqslant m$ parts, then H has an ε-homogeneous partition with $\leqslant 2^{\mathrm{poly}(m/\varepsilon)}$ parts.

VC_1 dimension

Theorem (Gishboliner–Shapira–W.)

If every link of H has an ε-homogeneous partition with $\leqslant m$ parts, then H has an ε-homogeneous partition with $\leqslant 2^{\mathrm{poly}(m/\varepsilon)}$ parts.

VC_1 dimension

Theorem (Gishboliner–Shapira–W.)

If every link of H has an ε-homogeneous partition with $\leqslant m$ parts, then H has an ε-homogeneous partition with $\leqslant 2^{\mathrm{poly}(m/\varepsilon)}$ parts.

We know which graphs have small homogeneous partitions!

VC_1 dimension

Theorem (Gishboliner–Shapira–W.)

If every link of H has an ε-homogeneous partition with $\leqslant m$ parts, then H has an ε-homogeneous partition with $\leqslant 2^{\mathrm{poly}(m/\varepsilon)}$ parts.

We know which graphs have small homogeneous partitions!

Definition

H has bounded VC_1 dimension if all links have bounded VC dim.

VC_1 dimension

Theorem (Gishboliner–Shapira–W.)

If every link of H has an ε-homogeneous partition with $\leqslant m$ parts, then H has an ε-homogeneous partition with $\leqslant 2^{\mathrm{poly}(m/\varepsilon)}$ parts.

We know which graphs have small homogeneous partitions!

Definition

H has bounded VC_1 dimension if all links have bounded VC dim.

Theorem (Chernikov–Towsner, Terry–Wolf)

If H has bounded VC_1 dimension, then H has an ε-homogeneous partition with $\leqslant M(\varepsilon)$ parts.

VC_1 dimension

Theorem (Gishboliner–Shapira–W.)

If every link of H has an ε-homogeneous partition with $\leqslant m$ parts, then H has an ε-homogeneous partition with $\leqslant 2^{\mathrm{poly}(m/\varepsilon)}$ parts.

We know which graphs have small homogeneous partitions!

Definition

H has bounded VC_1 dimension if all links have bounded VC dim.

Theorem (Chernikov–Towsner, Terry–Wolf, Terry)

If H has bounded VC_1 dimension, then H has an ε-homogeneous partition with $\leqslant M(\varepsilon)$ parts. Moreover, $2^{\mathrm{poly}(1/\varepsilon)} \leqslant M(\varepsilon) \leqslant 2^{2^{\mathrm{poly}(1/\varepsilon)}}$.

VC_1 dimension

Theorem (Gishboliner–Shapira–W.)

If every link of H has an ε-homogeneous partition with $\leqslant m$ parts, then H has an ε-homogeneous partition with $\leqslant 2^{\text{poly}(m/\varepsilon)}$ parts.

We know which graphs have small homogeneous partitions!

Definition

H has bounded VC_1 dimension if all links have bounded VC dim.

Theorem (Chernikov–Towsner, Terry–Wolf, Terry)

If H has bounded VC_1 dimension, then H has an ε-homogeneous partition with $\leqslant M(\varepsilon)$ parts. Moreover, $2^{\text{poly}(1/\varepsilon)} \leqslant M(\varepsilon) \leqslant 2^{2^{\text{poly}(1/\varepsilon)}}$.

Terry conjectured that double exponential is necessary in general.

VC_1 dimension

Theorem (Gishboliner–Shapira–W.)

If every link of H has an ε-homogeneous partition with $\leqslant m$ parts, then H has an ε-homogeneous partition with $\leqslant 2^{\mathrm{poly}(m/\varepsilon)}$ parts.

We know which graphs have small homogeneous partitions!

Definition

H has bounded VC_1 dimension if all links have bounded VC dim.

Theorem (Chernikov–Towsner, Terry–Wolf, Terry)

If H has bounded VC_1 dimension, then H has an ε-homogeneous partition with $\leqslant M(\varepsilon)$ parts. Moreover, $2^{\mathrm{poly}(1/\varepsilon)} \leqslant M(\varepsilon) \leqslant 2^{2^{\mathrm{poly}(1/\varepsilon)}}$.

Terry conjectured that double exponential is necessary in general.

Corollary (Gishboliner–Shapira–W.)

$M(\varepsilon) = 2^{\mathrm{poly}(1/\varepsilon)}$, *i.e., single exponential is necessary and sufficient.*

More characterizations

More characterizations

Theorem (Terry)

The following are equivalent for a hereditary family of 3-graphs H.

- *H has bounded strong VC dimension.*
- *H has an ε-homogeneous partition with $\leqslant \text{poly}(\frac{1}{\varepsilon})$ parts.*
- *H has a weak ε-regular partition with $\leqslant \text{poly}(\frac{1}{\varepsilon})$ parts.*

More characterizations

Theorem (Terry)

The following are equivalent for a hereditary family of 3-graphs H.

- *H has bounded strong VC dimension.*
- *H has an ε-homogeneous partition with* $\leqslant \operatorname{poly}(\frac{1}{\varepsilon})$ *parts.*
- *H has a weak ε-regular partition with* $\leqslant \operatorname{poly}(\frac{1}{\varepsilon})$ *parts.*

Bounded strong VC dimension implies bounded VC_1 dimension.

More characterizations

Theorem (Terry)

The following are equivalent for a hereditary family of 3-graphs H.

- *H has bounded strong VC dimension.*
- *H has an ε-homogeneous partition with $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts.*
- *H has a weak ε-regular partition with $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts.*

Bounded strong VC dimension implies bounded VC_1 dimension.

Theorem (Terry, CT, TW, GSW)

The following are equivalent for a hereditary family of 3-graphs H:

More characterizations

Theorem (Terry)

The following are equivalent for a hereditary family of 3-graphs H.

- *H has bounded strong VC dimension.*
- *H has an ε-homogeneous partition with $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts.*
- *H has a weak ε-regular partition with $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts.*

Bounded strong VC dimension implies bounded VC_1 dimension.

Theorem (Terry, CT, TW, GSW)

The following are equivalent for a hereditary family of 3-graphs H:

- *H has bounded VC_1 dimension.*

More characterizations

Theorem (Terry)

The following are equivalent for a hereditary family of 3-graphs H.

- *H has bounded strong VC dimension.*
- *H has an ε-homogeneous partition with* $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ *parts.*
- *H has a weak ε-regular partition with* $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ *parts.*

Bounded strong VC dimension implies bounded VC_1 dimension.

Theorem (Terry, CT, TW, GSW)

The following are equivalent for a hereditary family of 3-graphs H:

- *H has bounded* VC_1 *dimension.*
- *H has an ε-homogeneous partition with* $\leqslant 2^{\mathrm{poly}(1/\varepsilon)}$ *parts.*

More characterizations

Theorem (Terry)

The following are equivalent for a hereditary family of 3-graphs H.

- *H has bounded strong VC dimension.*
- *H has an ε-homogeneous partition with $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts.*
- *H has a weak ε-regular partition with $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts.*

Bounded strong VC dimension implies bounded VC_1 dimension.

Theorem (Terry, CT, TW, GSW)

The following are equivalent for a hereditary family of 3-graphs H:

- *H has bounded VC_1 dimension.*
- *H has an ε-homogeneous partition with $\leqslant 2^{\mathrm{poly}(1/\varepsilon)}$ parts.*
- *H has an ε-homogeneous partition with any number of parts.*

More characterizations

Theorem (Terry)

The following are equivalent for a hereditary family of 3-graphs H.

- *H has bounded strong VC dimension.*
- *H has an ε-homogeneous partition with $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts.*
- *H has a weak ε-regular partition with $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ parts.*

Bounded strong VC dimension implies bounded VC_1 dimension.

Theorem (Terry, CT, TW, GSW)

The following are equivalent for a hereditary family of 3-graphs H:

- *H has bounded VC_1 dimension.*
- *H has an ε-homogeneous partition with $\leqslant 2^{\mathrm{poly}(1/\varepsilon)}$ parts.*
- *H has an ε-homogeneous partition with any number of parts.*
- *H has a weak ε-regular partition with $\leqslant 2^{\mathrm{poly}(1/\varepsilon)}$ parts.*

More characterizations

Theorem (Terry)

The following are equivalent for a hereditary family of 3-graphs H.

- *H has bounded strong VC dimension.*
- *H has an ε-homogeneous partition with* $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ *parts.*
- *H has a weak ε-regular partition with* $\leqslant \mathrm{poly}(\frac{1}{\varepsilon})$ *parts.*

Bounded strong VC dimension implies bounded VC_1 dimension.

Theorem (Terry, CT, TW, GSW)

The following are equivalent for a hereditary family of 3-graphs H:

- *H has bounded* VC_1 *dimension.*
- *H has an ε-homogeneous partition with* $\leqslant 2^{\mathrm{poly}(1/\varepsilon)}$ *parts.*
- *H has an ε-homogeneous partition with any number of parts.*
- *H has a weak ε-regular partition with* $\leqslant 2^{\mathrm{poly}(1/\varepsilon)}$ *parts.*
- *H has a weak ε-regular partition with sub-tower number of parts.*

VC_2 dimension

Question

What does it mean for a 3-graph to be simple?

Question

What does it mean for a 3-graph to be simple?

Hypergraphs are scary! Let's try to only think about graphs.

VC_2 dimension

Question

What does it mean for a 3-graph to be simple?

Hypergraphs are scary! Let's try to only think about graphs.

G has bounded VC dimension if it forbids some fixed bi-induced B.

VC_2 dimension

Question

What does it mean for a 3-graph to be simple?

Hypergraphs are scary! Let's try to only think about graphs.

G has bounded VC dimension if it forbids some fixed bi-induced B.

The most natural extension to 3-graphs is:

Definition (Shelah)

H has bounded VC_2 dimension if it forbids some fixed tri-induced T.

Tri-induced: Edges across all 3 parts are induced, all other edges 🤷.

VC_2 dimension

Question

What does it mean for a 3-graph to be simple?

Hypergraphs are scary! Let's try to only think about graphs.

G has bounded VC dimension if it forbids some fixed bi-induced B.

The most natural extension to 3-graphs is:

Definition (Shelah)

H has bounded VC_2 dimension if it forbids some fixed tri-induced T.

Tri-induced: Edges across all 3 parts are induced, all other edges 🤷.

This is a natural combinatorial notion, but it originates in logic.

VC_2 dimension

Question

What does it mean for a 3-graph to be simple?

Hypergraphs are scary! Let's try to only think about graphs.

G has bounded VC dimension if it forbids some fixed bi-induced B.

The most natural extension to 3-graphs is:

Definition (Shelah)

H has bounded VC_2 dimension if it forbids some fixed tri-induced T.

Tri-induced: Edges across all 3 parts are induced, all other edges 🤷.

This is a natural combinatorial notion, but it originates in logic.

Bounded VC_1 dimension $\Longleftrightarrow$ forbidden bi-induced B in every link

VC_2 dimension

Question

What does it mean for a 3-graph to be simple?

Hypergraphs are scary! Let's try to only think about graphs.

G has bounded VC dimension if it forbids some fixed bi-induced B.

The most natural extension to 3-graphs is:

Definition (Shelah)

H has bounded VC_2 dimension if it forbids some fixed tri-induced T.

Tri-induced: Edges across all 3 parts are induced, all other edges 🤷.

This is a natural combinatorial notion, but it originates in logic.

Bounded VC_1 dimension $\iff$ forbidden bi-induced B in every link $\iff$ forbidden tri-induced T, where one part of T is a singleton.

VC_2 dimension

Question

What does it mean for a 3-graph to be simple?

Hypergraphs are scary! Let's try to only think about graphs.

G has bounded VC dimension if it forbids some fixed bi-induced B.

The most natural extension to 3-graphs is:

Definition (Shelah)

H has bounded VC_2 dimension if it forbids some fixed tri-induced T.

Tri-induced: Edges across all 3 parts are induced, all other edges 🤷.

This is a natural combinatorial notion, but it originates in logic.

Bounded VC_1 dimension $\iff$ forbidden bi-induced B in every link $\iff$ forbidden tri-induced T, where one part of T is a singleton.

$$\text{bounded } VC_1 \implies \text{bounded } VC_2.$$

VC_2 dimension

Question

What does it mean for a 3-graph to be simple?

Hypergraphs are scary! Let's try to only think about graphs.

G has bounded VC dimension if it forbids some fixed bi-induced B.

The most natural extension to 3-graphs is:

Definition (Shelah)

H has bounded VC_2 dimension if it forbids some fixed tri-induced T.

Tri-induced: Edges across all 3 parts are induced, all other edges 🤷.

This is a natural combinatorial notion, but it originates in logic.

Bounded VC_1 dimension $\iff$ forbidden bi-induced B in every link $\iff$ forbidden tri-induced T, where one part of T is a singleton.

Bounded strong VC $\implies$ bounded VC_1 $\implies$ bounded VC_2.

Example of bounded VC_2 dimension

H has bounded VC_2 dimension if it forbids some fixed tri-induced T.

Example of bounded VC_2 dimension

H has bounded VC_2 dimension if it forbids some fixed tri-induced T.

Example: Let G_1, G_2, G_3 be arbitrary bipartite graphs.

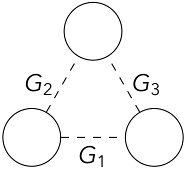

Example of bounded VC_2 dimension

H has bounded VC_2 dimension if it forbids some fixed tri-induced T.

Example: Let G_1, G_2, G_3 be arbitrary bipartite graphs. Let H be the hypergraph of all triangles in $G_1 \cup G_2 \cup G_3$, denoted $\Delta(G_1 \cup G_2 \cup G_3)$.

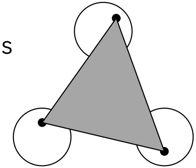

Example of bounded VC_2 dimension

H has bounded VC_2 dimension if it forbids some fixed tri-induced T.

Example: Let G_1, G_2, G_3 be arbitrary bipartite graphs. Let H be the hypergraph of all triangles in $G_1 \cup G_2 \cup G_3$, denoted $\Delta(G_1 \cup G_2 \cup G_3)$.

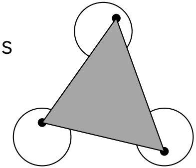

Claim: H forbids tri-induced $K^{(3)}_{2,2,2} \setminus e$

Example of bounded VC_2 dimension

H has bounded VC_2 dimension if it forbids some fixed tri-induced T.

Example: Let G_1, G_2, G_3 be arbitrary bipartite graphs. Let H be the hypergraph of all triangles in $G_1 \cup G_2 \cup G_3$, denoted $\Delta(G_1 \cup G_2 \cup G_3)$.

Claim: H forbids tri-induced $K^{(3)}_{2,2,2} \setminus e$

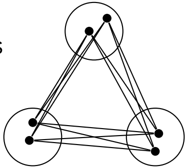

Example of bounded VC_2 dimension

H has bounded VC_2 dimension if it forbids some fixed tri-induced T.

Example: Let G_1, G_2, G_3 be arbitrary bipartite graphs. Let H be the hypergraph of all triangles in $G_1 \cup G_2 \cup G_3$, denoted $\Delta(G_1 \cup G_2 \cup G_3)$.

Claim: H forbids tri-induced $K_{2,2,2}^{(3)} \setminus e$, hence has bounded VC_2 dimension.

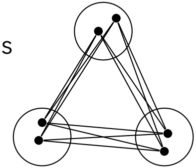

Example of bounded VC_2 dimension

H has bounded VC_2 dimension if it forbids some fixed tri-induced *T*.

Example: Let G_1, G_2, G_3 be arbitrary bipartite graphs. Let H be the hypergraph of all triangles in $G_1 \cup G_2 \cup G_3$, denoted $\Delta(G_1 \cup G_2 \cup G_3)$.

Claim: H forbids tri-induced $K^{(3)}_{2,2,2} \setminus e$, hence has bounded VC_2 dimension.

This example is very versatile.

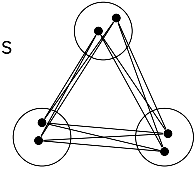

Example of bounded VC_2 dimension

H has bounded VC_2 dimension if it forbids some fixed tri-induced T.

Example: Let G_1, G_2, G_3 be arbitrary bipartite graphs. Let H be the hypergraph of all triangles in $G_1 \cup G_2 \cup G_3$, denoted $\Delta(G_1 \cup G_2 \cup G_3)$.

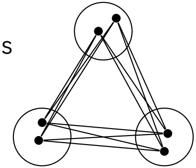

Claim: H forbids tri-induced $K^{(3)}_{2,2,2} \setminus e$, hence has bounded VC_2 dimension.

This example is very versatile.

H can have bounded VC_2 dimension, yet require tower-type bounds in Chung's regularity lemma.

Example of bounded VC_2 dimension

H has bounded VC_2 dimension if it forbids some fixed tri-induced T.

Example: Let G_1, G_2, G_3 be arbitrary bipartite graphs. Let H be the hypergraph of all triangles in $G_1 \cup G_2 \cup G_3$, denoted $\Delta(G_1 \cup G_2 \cup G_3)$.

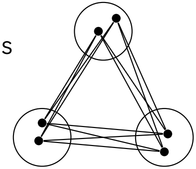

Claim: H forbids tri-induced $K^{(3)}_{2,2,2} \setminus e$, hence has bounded VC_2 dimension.

This example is very versatile.

H can have bounded VC_2 dimension, yet require tower-type bounds in Chung's regularity lemma.

Proof: Take G_i to be Gowers's graphs.

Example of bounded VC_2 dimension

H has bounded VC_2 dimension if it forbids some fixed tri-induced T.

Example: Let G_1, G_2, G_3 be arbitrary bipartite graphs. Let H be the hypergraph of all triangles in $G_1 \cup G_2 \cup G_3$, denoted $\Delta(G_1 \cup G_2 \cup G_3)$.

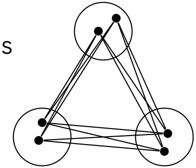

Claim: H forbids tri-induced $K^{(3)}_{2,2,2} \setminus e$, hence has bounded VC_2 dimension.

This example is very versatile.

H can have bounded VC_2 dimension, yet require tower-type bounds in Chung's regularity lemma.

Proof: Take G_i to be Gowers's graphs.

H can have bounded VC_2 dimension, yet not admit any ε-homogeneous partition of bounded size.

Example of bounded VC_2 dimension

H has bounded VC_2 dimension if it forbids some fixed tri-induced T.

Example: Let G_1, G_2, G_3 be arbitrary bipartite graphs. Let H be the hypergraph of all triangles in $G_1 \cup G_2 \cup G_3$, denoted $\Delta(G_1 \cup G_2 \cup G_3)$.

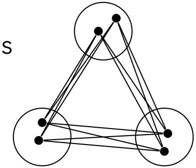

Claim: H forbids tri-induced $K^{(3)}_{2,2,2} \setminus e$, hence has bounded VC_2 dimension.

This example is very versatile.

H can have bounded VC_2 dimension, yet require tower-type bounds in Chung's regularity lemma.

Proof: Take G_i to be Gowers's graphs.

H can have bounded VC_2 dimension, yet not admit any ε-homogeneous partition of bounded size.

Proof: Take G_i to be random of density $\frac{1}{2}$.

Example of bounded VC_2 dimension

H has bounded VC_2 dimension if it forbids some fixed tri-induced T.

Example: Let G_1, G_2, G_3 be arbitrary bipartite graphs. Let H be the hypergraph of all triangles in $G_1 \cup G_2 \cup G_3$, denoted $\Delta(G_1 \cup G_2 \cup G_3)$.

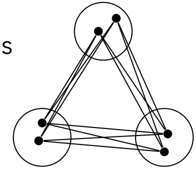

Claim: H forbids tri-induced $K_{2,2,2}^{(3)} \setminus e$, hence has bounded VC_2 dimension.

This example is very versatile.

H can have bounded VC_2 dimension, yet require tower-type bounds in Chung's regularity lemma.

Proof: Take G_i to be Gowers's graphs.

H can have bounded VC_2 dimension, yet not admit any ε-homogeneous partition of bounded size.

Proof: Take G_i to be random of density $\frac{1}{2}$.

Weak regularity does not support a(n induced) counting lemma.

Example of bounded VC_2 dimension

H has bounded VC_2 dimension if it forbids some fixed tri-induced T.

Example: Let G_1, G_2, G_3 be arbitrary bipartite graphs. Let H be the hypergraph of all triangles in $G_1 \cup G_2 \cup G_3$, denoted $\Delta(G_1 \cup G_2 \cup G_3)$.

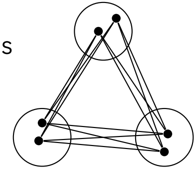

Claim: H forbids tri-induced $K^{(3)}_{2,2,2} \setminus e$, hence has bounded VC_2 dimension.

This example is very versatile.

H can have bounded VC_2 dimension, yet require tower-type bounds in Chung's regularity lemma.

Proof: Take G_i to be Gowers's graphs.

H can have bounded VC_2 dimension, yet not admit any ε-homogeneous partition of bounded size.

Proof: Take G_i to be random of density $\frac{1}{2}$.

Weak regularity does not support a(n induced) counting lemma.

Proof: Take G_i to be random of density $\frac{1}{2}$.

The full hypergraph regularity lemma

Weak regularity does not support a(n induced) counting lemma.

Proof: Take G_i to be random of density $\frac{1}{2}$.

The full hypergraph regularity lemma

Weak regularity does not support a(n induced) counting lemma.

Proof: Take G_i to be random of density $\frac{1}{2}$.

This example explains the weakness of weak regularity, but also suggests what goes wrong:

The full hypergraph regularity lemma

Weak regularity does not support a(n induced) counting lemma.

Proof: Take G_i to be random of density $\frac{1}{2}$.

This example explains the weakness of weak regularity, but also suggests what goes wrong: a hypergraph can depend on "lower-uniformity information", which must be taken into account.

The full hypergraph regularity lemma

Weak regularity does not support a(n induced) counting lemma.

Proof: Take G_i to be random of density $\frac{1}{2}$.

This example explains the weakness of weak regularity, but also suggests what goes wrong: a hypergraph can depend on "lower-uniformity information", which must be taken into account.

A hypergraph regularity partition of H consists of:

The full hypergraph regularity lemma

Weak regularity does not support a(n induced) counting lemma.

Proof: Take G_i to be random of density $\frac{1}{2}$.

This example explains the weakness of weak regularity, but also suggests what goes wrong: a hypergraph can depend on "lower-uniformity information", which must be taken into account.

A hypergraph regularity partition of H consists of:

- a partition $V(H) = V_1 \sqcup \cdots \sqcup V_m$,

The full hypergraph regularity lemma

Weak regularity does not support a(n induced) counting lemma.

Proof: Take G_i to be random of density $\frac{1}{2}$.

This example explains the weakness of weak regularity, but also suggests what goes wrong: a hypergraph can depend on "lower-uniformity information", which must be taken into account.

A hypergraph regularity partition of H consists of:

- a partition $V(H) = V_1 \sqcup \cdots \sqcup V_m$,
- a partition of each $V_i \times V_j$ into bipartite graphs $G_{ij}^1, \ldots, G_{ij}^t$,

The full hypergraph regularity lemma

Weak regularity does not support a(n induced) counting lemma.

Proof: Take G_i to be random of density $\frac{1}{2}$.

This example explains the weakness of weak regularity, but also suggests what goes wrong: a hypergraph can depend on "lower-uniformity information", which must be taken into account.

A hypergraph regularity partition of H consists of:

- a partition $V(H) = V_1 \sqcup \cdots \sqcup V_m$,
- a partition of each $V_i \times V_j$ into bipartite graphs $G_{ij}^1, \ldots, G_{ij}^t$,
- s.t. $E(H)$ "lies ε-regularly in" $\Delta(G_{ij}^a \cup G_{jk}^b \cup G_{ik}^c)$ for most triples,

The full hypergraph regularity lemma

Weak regularity does not support a(n induced) counting lemma.

Proof: Take G_i to be random of density $\frac{1}{2}$.

This example explains the weakness of weak regularity, but also suggests what goes wrong: a hypergraph can depend on "lower-uniformity information", which must be taken into account.

A hypergraph regularity partition of H consists of:

- a partition $V(H) = V_1 \sqcup \cdots \sqcup V_m$,
- a partition of each $V_i \times V_j$ into bipartite graphs $G_{ij}^1, \ldots, G_{ij}^t$,
- s.t. $E(H)$ "lies ε-regularly in" $\Delta(G_{ij}^a \cup G_{jk}^b \cup G_{ik}^c)$ for most triples,
- and each G_{ij}^a is very regular.

The full hypergraph regularity lemma

Weak regularity does not support a(n induced) counting lemma.

Proof: Take G_i to be random of density $\frac{1}{2}$.

This example explains the weakness of weak regularity, but also suggests what goes wrong: a hypergraph can depend on "lower-uniformity information", which must be taken into account.

A hypergraph regularity partition of H consists of:

- a partition $V(H) = V_1 \sqcup \cdots \sqcup V_m$,
- a partition of each $V_i \times V_j$ into bipartite graphs $G_{ij}^1, \ldots, G_{ij}^t$,
- s.t. $E(H)$ "lies ε-regularly in" $\Delta(G_{ij}^a \cup G_{jk}^b \cup G_{ik}^c)$ for most triples,
- and each G_{ij}^a is very regular.

Very: G_{ij}^a has edge density $\frac{1}{t}$, so we require $f(\frac{1}{t})$-regularity.

The full hypergraph regularity lemma

Weak regularity does not support a(n induced) counting lemma.

Proof: Take G_i to be random of density $\frac{1}{2}$.

This example explains the weakness of weak regularity, but also suggests what goes wrong: a hypergraph can depend on "lower-uniformity information", which must be taken into account.

A hypergraph regularity partition of H consists of:

- a partition $V(H) = V_1 \sqcup \cdots \sqcup V_m$,
- a partition of each $V_i \times V_j$ into bipartite graphs $G_{ij}^1, \ldots, G_{ij}^t$,
- s.t. $E(H)$ "lies ε-regularly in" $\Delta(G_{ij}^a \cup G_{jk}^b \cup G_{ik}^c)$ for most triples,
- and each G_{ij}^a is very regular.

Very: G_{ij}^a has edge density $\frac{1}{t}$, so we require $f(\frac{1}{t})$-regularity.

Theorem (Hypergraph regularity lemma)

Every H has an (ε, f)-regular partition with $\leqslant M_f(\varepsilon)$ parts.

Hypergraph regularity and VC_2 dimension

A hypergraph regularity partition of H consists of:

- a partition $V(H) = V_1 \sqcup \cdots \sqcup V_m$,
- a partition of each $V_i \times V_j$ into bipartite graphs $G_{ij}^1, \ldots, G_{ij}^t$,
- s.t. $E(H)$ "lies ε-regularly in" $\Delta(G_{ij}^a \cup G_{jk}^b \cup G_{ik}^c)$ for most triples,
- and each G_{ij}^a is very regular, i.e. $f(\frac{1}{t})$-regular.

Theorem (Hypergraph regularity lemma)

Every H has an (ε, f)-regular partition with $\leqslant M_f(\varepsilon)$ parts.

Hypergraph regularity and VC_2 dimension

A hypergraph regularity partition of H consists of:

- a partition $V(H) = V_1 \sqcup \cdots \sqcup V_m$,
- a partition of each $V_i \times V_j$ into bipartite graphs $G_{ij}^1, \ldots, G_{ij}^t$,
- s.t. $E(H)$ "lies ε-regularly in" $\Delta(G_{ij}^a \cup G_{jk}^b \cup G_{ik}^c)$ for most triples,
- and each G_{ij}^a is very regular, i.e. $f(\frac{1}{t})$-regular.

Theorem (Hypergraph regularity lemma)

Every H has an (ε, f)-regular partition with $\leqslant M_f(\varepsilon)$ parts.

Theorem (Chernikov–Towsner, Terry–Wolf)

If H has bounded VC_2 dimension, then $E(H)$ lies ε-homogeneously in $\Delta(G_{ij}^a \cup G_{jk}^b \cup G_{ik}^c)$.

Hypergraph regularity and VC_2 dimension

A hypergraph regularity partition of H consists of:

- a partition $V(H) = V_1 \sqcup \cdots \sqcup V_m$,
- a partition of each $V_i \times V_j$ into bipartite graphs $G_{ij}^1, \ldots, G_{ij}^t$,
- s.t. $E(H)$ "lies ε-regularly in" $\Delta(G_{ij}^a \cup G_{jk}^b \cup G_{ik}^c)$ for most triples,
- and each G_{ij}^a is very regular, i.e. $f(\frac{1}{t})$-regular.

Theorem (Hypergraph regularity lemma)

Every H has an (ε, f)-regular partition with $\leqslant M_f(\varepsilon)$ parts.

Theorem (Chernikov–Towsner, Terry–Wolf)

If H has bounded VC_2 dimension, then $E(H)$ lies ε-homogeneously in $\Delta(G_{ij}^a \cup G_{jk}^b \cup G_{ik}^c)$.

If H has bounded VC_2 dimension, all the information of H comes from uniformity 2 $\Longleftrightarrow$ the example we saw is the only example.

Quantitative aspects of hypergraph regularity

Theorem (Hypergraph regularity lemma)

Every H has an (ε, f)-regular partition with $\leqslant M_f(\varepsilon)$ parts.

Quantitative aspects of hypergraph regularity

Theorem (Hypergraph regularity lemma)

Every H has an (ε, f)-regular partition with $\leqslant M_f(\varepsilon)$ parts.

The proof iterates Szemerédi's regularity lemma, so gives

$$M_f(\varepsilon) \leqslant \text{wowzer}(\varepsilon^{-C}) := \underbrace{2^{2^{\cdot^{\cdot^{\cdot^{2}}}}}\Big\} 2^{2^{\cdot^{\cdot^{\cdot^{2}}}}}\Big\} \cdots 2}_{\varepsilon^{-C}\text{ times}}$$

Quantitative aspects of hypergraph regularity

Theorem (Hypergraph regularity lemma)

Every H has an (ε, f)-regular partition with $\leqslant M_f(\varepsilon)$ parts.

The proof iterates Szemerédi's regularity lemma, so gives

$$M_f(\varepsilon) \leqslant \text{wowzer}(\varepsilon^{-C}) := \underbrace{2^{2^{\cdot^{\cdot^{\cdot^{2}}}}} \Big\} 2^{2^{\cdot^{\cdot^{\cdot^{2}}}}} \Big\} \cdots 2}_{\varepsilon^{-C} \text{ times}}$$

for "sane" functions f.

Quantitative aspects of hypergraph regularity

Theorem (Hypergraph regularity lemma)

Every H has an (ε, f)-regular partition with $\leqslant M_f(\varepsilon)$ parts.

The proof iterates Szemerédi's regularity lemma, so gives

$$M_f(\varepsilon) \leqslant \text{wowzer}(\varepsilon^{-C}) := \underbrace{2^{2^{\cdot^{\cdot^{\cdot^{2}}}}} \Big\} 2^{2^{\cdot^{\cdot^{\cdot^{2}}}}} \Big\} \cdots 2}_{\varepsilon^{-C} \text{ times}}$$

for "sane" functions f. If f is insane, the bounds are even worse.

Quantitative aspects of hypergraph regularity

Theorem (Hypergraph regularity lemma)

Every H has an (ε, f)-regular partition with $\leqslant M_f(\varepsilon)$ parts.

The proof iterates Szemerédi's regularity lemma, so gives

$$M_f(\varepsilon) \leqslant \text{wowzer}(\varepsilon^{-C}) := \underbrace{2^{2^{\cdot^{\cdot^{\cdot^{2}}}}} \Big\} 2^{2^{\cdot^{\cdot^{\cdot^{2}}}}} \Big\} \cdots 2}_{\varepsilon^{-C} \text{ times}}$$

for "sane" functions f. If f is insane, the bounds are even worse.

Most applications, e.g. the hypergraph removal lemma, use $f(x) = \text{poly}(x)$, which is very sane.

Quantitative aspects of hypergraph regularity

Theorem (Hypergraph regularity lemma)

Every H has an (ε, f)-regular partition with $\leqslant M_f(\varepsilon)$ parts.

The proof iterates Szemerédi's regularity lemma, so gives

$$M_f(\varepsilon) \leqslant \text{wowzer}(\varepsilon^{-C}) := \underbrace{2^{2^{\cdot^{\cdot^{\cdot^{2}}}}} \Big\} 2^{2^{\cdot^{\cdot^{\cdot^{2}}}}} \Big\} \cdots 2}_{\varepsilon^{-C} \text{ times}}$$

for "sane" functions f. If f is insane, the bounds are even worse.

Most applications, e.g. the hypergraph removal lemma, use $f(x) = \text{poly}(x)$, which is very sane.

Theorem (Moshkovitz–Shapira)

Wowzer-type bounds are necessary, even for very sane f.

Quantitative aspects under bounded VC_2 dimension

Theorem (Hypergraph regularity lemma)

Every H has an (ε, f)-regular partition with $\leqslant M_f(\varepsilon)$ parts.

Theorem (Moshkovitz–Shapira)

Wowzer-type bounds are necessary, even for very sane f.

Quantitative aspects under bounded VC_2 dimension

Theorem (Hypergraph regularity lemma)

Every H has an (ε, f)-regular partition with $\leqslant M_f(\varepsilon)$ parts.

Theorem (Moshkovitz–Shapira)

Wowzer-type bounds are necessary, even for very sane f.

Theorem (Gishboliner–Shapira–W.)

If H has bounded VC_2 dimension, $M_f(\varepsilon) \leqslant \mathrm{twr}(\varepsilon^{-C})$ for sane f (and this is best possible).

Quantitative aspects under bounded VC_2 dimension

Theorem (Hypergraph regularity lemma)

Every H has an (ε, f)-regular partition with $\leqslant M_f(\varepsilon)$ parts.

Theorem (Moshkovitz–Shapira)

Wowzer-type bounds are necessary, even for very sane f.

Theorem (Gishboliner–Shapira–W.)

If H has bounded VC_2 dimension, $M_f(\varepsilon) \leqslant \mathrm{twr}(\varepsilon^{-C})$ for sane f (and this is best possible).

Upshot: 3-graphs of bounded VC_2 dimension act like graphs, even with respect to the bounds in the regularity lemma.

Quantitative aspects under bounded VC_2 dimension

Theorem (Hypergraph regularity lemma)

Every H has an (ε, f)-regular partition with $\leqslant M_f(\varepsilon)$ parts.

Theorem (Moshkovitz–Shapira)

Wowzer-type bounds are necessary, even for very sane f.

Theorem (Gishboliner–Shapira–W.)

If H has bounded VC_2 dimension, $M_f(\varepsilon) \leqslant \operatorname{twr}(\varepsilon^{-C})$ for sane f (and this is best possible).

Upshot: 3-graphs of bounded VC_2 dimension act like graphs, even with respect to the bounds in the regularity lemma.

Theorem (Terry)

If H has bounded VC_2 dimension and f is arbitrary, then wowzer-type bounds are necessary.

Higher uniformities

Higher uniformities

In a hypergraph regularity partition of a k-graph, we have to partition the vertices, pairs, triples, ..., $(k-1)$-tuples.

Higher uniformities

In a hypergraph regularity partition of a k-graph, we have to partition the vertices, pairs, triples, ..., $(k-1)$-tuples.

The proof iterates the $(k-1)$-uniform regularity lemma, so gives bounds on the kth level of the Ackermann hierarchy.

Higher uniformities

In a hypergraph regularity partition of a k-graph, we have to partition the vertices, pairs, triples, ..., $(k-1)$-tuples.

The proof iterates the $(k-1)$-uniform regularity lemma, so gives bounds on the kth level of the Ackermann hierarchy.

Moshkovitz–Shapira: Ackermann$_k$ bounds are necessary in general.

Higher uniformities

In a hypergraph regularity partition of a k-graph, we have to partition the vertices, pairs, triples, …, $(k-1)$-tuples.

The proof iterates the $(k-1)$-uniform regularity lemma, so gives bounds on the kth level of the Ackermann hierarchy.

Moshkovitz–Shapira: Ackermann$_k$ bounds are necessary in general.

A k-graph H has bounded VC$_{k-1}$ dimension if it forbids a fixed k-induced k-partite k-graph.

Higher uniformities

In a hypergraph regularity partition of a k-graph, we have to partition the vertices, pairs, triples, ..., $(k-1)$-tuples.

The proof iterates the $(k-1)$-uniform regularity lemma, so gives bounds on the kth level of the Ackermann hierarchy.

Moshkovitz–Shapira: Ackermann$_k$ bounds are necessary in general.

A k-graph H has bounded VC$_{k-1}$ dimension if it forbids a fixed k-induced k-partite k-graph. It has bounded VC$_r$ dimension ($r < k$) if all $(r+1)$-uniform links have bounded VC$_r$ dimension.

Higher uniformities

In a hypergraph regularity partition of a k-graph, we have to partition the vertices, pairs, triples, ..., $(k-1)$-tuples.

The proof iterates the $(k-1)$-uniform regularity lemma, so gives bounds on the kth level of the Ackermann hierarchy.

Moshkovitz–Shapira: Ackermann$_k$ bounds are necessary in general.

A k-graph H has bounded VC$_{k-1}$ dimension if it forbids a fixed k-induced k-partite k-graph. It has bounded VC$_r$ dimension ($r < k$) if all $(r+1)$-uniform links have bounded VC$_r$ dimension.

Theorem (Chernikov–Towsner)

If H has bounded VC$_r$ dimension, then it has a regularity partition that is homogeneous at uniformity $> r$.

Higher uniformities

In a hypergraph regularity partition of a k-graph, we have to partition the vertices, pairs, triples, ..., $(k-1)$-tuples.

The proof iterates the $(k-1)$-uniform regularity lemma, so gives bounds on the kth level of the Ackermann hierarchy.

Moshkovitz–Shapira: Ackermann$_k$ bounds are necessary in general.

A k-graph H has bounded VC$_{k-1}$ dimension if it forbids a fixed k-induced k-partite k-graph. It has bounded VC$_r$ dimension ($r < k$) if all $(r+1)$-uniform links have bounded VC$_r$ dimension.

Theorem (Chernikov–Towsner)

If H has bounded VC$_r$ dimension, then it has a regularity partition that is homogeneous at uniformity $> r$.

Theorem (Gishboliner–Shapira–W.)

If H has bounded VC$_r$ dimension, then it has a regularity partition with only Ackermann$_r$ *parts (and this is best possible).*

Higher uniformities

In a hypergraph regularity partition of a k-graph, we have to partition the vertices, pairs, triples, ..., $(k-1)$-tuples.

The proof iterates the $(k-1)$-uniform regularity lemma, so gives bounds on the kth level of the Ackermann hierarchy.

Moshkovitz–Shapira: Ackermann_k bounds are necessary in general.

A k-graph H has bounded VC_{k-1} dimension if it forbids a fixed k-induced k-partite k-graph. It has bounded VC_r dimension ($r < k$) if all $(r+1)$-uniform links have bounded VC_r dimension.

Theorem (Chernikov–Towsner)

If H has bounded VC_r dimension, then it has a regularity partition that is homogeneous at uniformity $> r$.

Theorem (Gishboliner–Shapira–W.)

If H has bounded VC_r dimension, then it has a regularity partition with only Ackermann_r *parts (and this is best possible).*

Upshot: Bounded VC_r dimension $\iff$ “looks like an r-graph”.

Proof non-sketch

If every link of H has an ε-homogeneous partition with $\leqslant m$ parts, then H has an ε-homogeneous partition with $\leqslant 2^{\mathrm{poly}(m/\varepsilon)}$ parts.

If H has bounded VC_2 dimension, $M_f(\varepsilon) \leqslant \mathrm{twr}(\varepsilon^{-C})$ for sane f (and this is best possible).

If H has bounded VC_r dimension, then it has a regularity partition with only $\mathrm{Ackermann}_r$ parts (and this is best possible).

Proof non-sketch

If every link of H has an ε-homogeneous partition with $\leqslant m$ parts, then H has an ε-homogeneous partition with $\leqslant 2^{\mathrm{poly}(m/\varepsilon)}$ parts.

If H has bounded VC_2 dimension, $M_f(\varepsilon) \leqslant \mathrm{twr}(\varepsilon^{-C})$ for sane f (and this is best possible).

If H has bounded VC_r dimension, then it has a regularity partition with only $\mathrm{Ackermann}_r$ parts (and this is best possible).

Traditional proofs don't extend, because these VC_r notions don't support a version of Haussler's packing lemma.

Proof non-sketch

If every link of H has an ε-homogeneous partition with $\leqslant m$ parts, then H has an ε-homogeneous partition with $\leqslant 2^{\mathrm{poly}(m/\varepsilon)}$ parts.

If H has bounded VC_2 dimension, $M_f(\varepsilon) \leqslant \mathrm{twr}(\varepsilon^{-C})$ for sane f (and this is best possible).

If H has bounded VC_r dimension, then it has a regularity partition with only $\mathrm{Ackermann}_r$ parts (and this is best possible).

Traditional proofs don't extend, because these VC_r notions don't support a version of Haussler's packing lemma.

Unifying theme: "Common refinement" is the enemy.

Proof non-sketch

If every link of H has an ε-homogeneous partition with $\leqslant m$ parts, then H has an ε-homogeneous partition with $\leqslant 2^{\text{poly}(m/\varepsilon)}$ parts.

If H has bounded VC_2 dimension, $M_f(\varepsilon) \leqslant \text{twr}(\varepsilon^{-C})$ for sane f (and this is best possible).

If H has bounded VC_r dimension, then it has a regularity partition with only Ackermann_r parts (and this is best possible).

Traditional proofs don't extend, because these VC_r notions don't support a version of Haussler's packing lemma.

Unifying theme: "Common refinement" is the enemy. Work with unconventional partitions for as long as possible.

Proof non-sketch

If every link of H has an ε-homogeneous partition with $\leqslant m$ parts, then H has an ε-homogeneous partition with $\leqslant 2^{\mathrm{poly}(m/\varepsilon)}$ parts.

If H has bounded VC_2 dimension, $M_f(\varepsilon) \leqslant \mathrm{twr}(\varepsilon^{-C})$ for sane f (and this is best possible).

We prove a cylinder regularity lemma for hypergraphs.

If H has bounded VC_r dimension, then it has a regularity partition with only $\mathrm{Ackermann}_r$ parts (and this is best possible).

Traditional proofs don't extend, because these VC_r notions don't support a version of Haussler's packing lemma.

Unifying theme: "Common refinement" is the enemy. Work with unconventional partitions for as long as possible.

Proof non-sketch

If every link of H has an ε-homogeneous partition with $\leqslant m$ parts, then H has an ε-homogeneous partition with $\leqslant 2^{\text{poly}(m/\varepsilon)}$ parts.

If H has bounded VC_2 dimension, $M_f(\varepsilon) \leqslant \text{twr}(\varepsilon^{-C})$ for sane f (and this is best possible).

We prove a cylinder regularity lemma for hypergraphs.

If H has bounded VC_r dimension, then it has a regularity partition with only Ackermann_r parts (and this is best possible).

Cylinder regularity at multiple uniformities simultaneously.

Traditional proofs don't extend, because these VC_r notions don't support a version of Haussler's packing lemma.

Unifying theme: "Common refinement" is the enemy. Work with unconventional partitions for as long as possible.

Proof non-sketch

If every link of H has an ε-homogeneous partition with $\leqslant m$ parts, then H has an ε-homogeneous partition with $\leqslant 2^{\mathrm{poly}(m/\varepsilon)}$ parts.

Partitions of vertices $\times$ pairs.

If H has bounded VC_2 dimension, $M_f(\varepsilon) \leqslant \mathrm{twr}(\varepsilon^{-C})$ for sane f (and this is best possible).

We prove a cylinder regularity lemma for hypergraphs.

If H has bounded VC_r dimension, then it has a regularity partition with only $\mathrm{Ackermann}_r$ parts (and this is best possible).

Cylinder regularity at multiple uniformities simultaneously.

Traditional proofs don't extend, because these VC_r notions don't support a version of Haussler's packing lemma.

Unifying theme: "Common refinement" is the enemy. Work with unconventional partitions for as long as possible.

Conclusion

Conclusion

There are several natural notions of simple hypergraphs.

Conclusion

There are several natural notions of simple hypergraphs.

They interact naturally with regularity, and characterize when we can get qualitative/quantitative improvements on regularity lemmas.

Conclusion

There are several natural notions of simple hypergraphs.

They interact naturally with regularity, and characterize when we can get qualitative/quantitative improvements on regularity lemmas.

They control how and when a hypergraph "looks like" a lower-uniformity hypergraph.

Thank you!