Arithmetic regularity, removal, and progressions

Jacob Fox
Stanford University

Matematické Kolokvium 99

Prague

November 22, 2016

Roth's theorem

Roth's theorem

Theorem (Roth)

Every subset $A \subset [N]$ with no three-term arithmetic progression has $|A| = o(N)$.

Roth's theorem

Theorem (Roth)

Every subset $A \subset [N]$ with no three-term arithmetic progression has $|A| = o(N)$.

Roth: $|A| = O(N/\log\log N)$.

Roth's theorem

Theorem (Roth)

Every subset $A \subset [N]$ with no three-term arithmetic progression has $|A| = o(N)$.

Roth: $|A| = O(N/\log\log N)$.

Improvements by Heath-Brown, Szemerédi, Bourgain.

Theorem (Roth)

Every subset $A \subset [N]$ with no three-term arithmetic progression has $|A| = o(N)$.

Roth: $|A| = O(N/\log\log N)$.

Improvements by Heath-Brown, Szemerédi, Bourgain.

Best known: $|A| \leq N/(\log N)^{1-o(1)}$ by Sanders, Bloom.

Roth's theorem

Theorem (Roth)

Every subset $A \subset [N]$ with no three-term arithmetic progression has $|A| = o(N)$.

Roth: $|A| = O(N/\log\log N)$.

Improvements by Heath-Brown, Szemerédi, Bourgain.

Best known: $|A| \leq N/(\log N)^{1-o(1)}$ by Sanders, Bloom.

Behrend construction gives a lower bound of $\frac{N}{e^{c\sqrt{\log N}}}$.

Game of Set

Game of Set

81 cards corresponding to points in $\mathbb{F}_3^4$.

Game of Set

81 cards corresponding to points in $\mathbb{F}_3^4$.

Question

How many cards can we have without a "set"?

Game of Set

81 cards corresponding to points in $\mathbb{F}_3^4$.

Question

How many cards can we have without a “set”?

Answer: 20

Cap Set Problem

Question

How large can $A \subset \mathbb{F}_3^n$ be without a 3-term arithmetic progression?

Cap Set Problem

Question

How large can $A \subset \mathbb{F}_3^n$ be without a 3-term arithmetic progression?

This variant of Roth's theorem is related to several famous problems in combinatorics and computer science, including the matrix multiplication problem, the sunflower conjecture.

Cap Set Problem

Question

How large can $A \subset \mathbb{F}_3^n$ be without a 3-term arithmetic progression?

This variant of Roth's theorem is related to several famous problems in combinatorics and computer science, including the matrix multiplication problem, the sunflower conjecture.

Brown-Buhler: $|A| = o(N)$.

Cap Set Problem

Question

How large can $A \subset \mathbb{F}_3^n$ be without a 3-term arithmetic progression?

This variant of Roth's theorem is related to several famous problems in combinatorics and computer science, including the matrix multiplication problem, the sunflower conjecture.

Brown-Buhler: $|A| = o(N)$.

Meshulam: $|A| = O(3^n/n)$.

Cap Set Problem

Question

How large can $A \subset \mathbb{F}_3^n$ be without a 3-term arithmetic progression?

This variant of Roth's theorem is related to several famous problems in combinatorics and computer science, including the matrix multiplication problem, the sunflower conjecture.

Brown-Buhler: $|A| = o(N)$.

Meshulam: $|A| = O(3^n/n)$.

Bateman-Katz: $|A| = O(3^n/n^{1+c})$.

Breakthrough

Theorem (Croot, Lev, Pach)

If $A \subset \mathbb{Z}_4^n$ has no 3-AP, then $|A| \leq 4^{cn}$ with $c \approx .926$.

Breakthrough

Theorem (Croot, Lev, Pach)

If $A \subset \mathbb{Z}_4^n$ has no 3-AP, then $|A| \leq 4^{cn}$ with $c \approx .926$.

Theorem (Ellenberg, Gijswijt)

If $A \subset \mathbb{F}_p^n$ has no 3-AP, then $|A| \leq p^{(1-c_p)n}$ for an explicit $c_p > 0$.

Breakthrough

Theorem (Croot, Lev, Pach)

If $A \subset \mathbb{Z}_4^n$ has no 3-AP, then $|A| \leq 4^{cn}$ with $c \approx .926$.

Theorem (Ellenberg, Gijswijt)

If $A \subset \mathbb{F}_p^n$ has no 3-AP, then $|A| \leq p^{(1-c_p)n}$ for an explicit $c_p > 0$.

Blasiak-Church-Cohn-Grochow-Naslund-Sawin-Umans, Alon

Same conclusion for the *multicolored sum-free problem*:

Breakthrough

Theorem (Croot, Lev, Pach)

If $A \subset \mathbb{Z}_4^n$ has no 3-AP, then $|A| \leq 4^{cn}$ with $c \approx .926$.

Theorem (Ellenberg, Gijswijt)

If $A \subset \mathbb{F}_p^n$ has no 3-AP, then $|A| \leq p^{(1-c_p)n}$ for an explicit $c_p > 0$.

Blasiak-Church-Cohn-Grochow-Naslund-Sawin-Umans, Alon

Same conclusion for the *multicolored sum-free problem*: If $\{x_i\}_{i=1}^m, \{y_i\}_{i=1}^m, \{z_i\}_{i=1}^m \subset \mathbb{F}_p^n$ with $x_i + y_j + z_k = 0 \Leftrightarrow i = j = k$, then $m \leq p^{(1-c_p)n}$.

Breakthrough

Theorem (Croot, Lev, Pach)

If $A \subset \mathbb{Z}_4^n$ has no 3-AP, then $|A| \leq 4^{cn}$ with $c \approx .926$.

Theorem (Ellenberg, Gijswijt)

If $A \subset \mathbb{F}_p^n$ has no 3-AP, then $|A| \leq p^{(1-c_p)n}$ for an explicit $c_p > 0$.

Blasiak-Church-Cohn-Grochow-Naslund-Sawin-Umans, Alon

Same conclusion for the *multicolored sum-free problem*: If $\{x_i\}_{i=1}^m, \{y_i\}_{i=1}^m, \{z_i\}_{i=1}^m \subset \mathbb{F}_p^n$ with $x_i + y_j + z_k = 0 \Leftrightarrow i = j = k$, then $m \leq p^{(1-c_p)n}$.

Theorem

Exponent is sharp for the *multicolored sum-free problem*: for $\mathbb{F}_2$ by construction of Fu-Kleinberg, $\mathbb{F}_p$ by Kleinberg-Sawin-Speyer.

Proof sketch of the multicolored sum-free problem

Slice rank of tensors: Tao

Proof sketch of the multicolored sum-free problem

Slice rank of tensors: Tao

A tensor $T : [N]^3 \to \mathbb{F}$ has slice rank 1 if there are functions $f : [N] \to \mathbb{F}$ and $g : [N]^2 \to \mathbb{F}$ such that one of the following holds:

Proof sketch of the multicolored sum-free problem

Slice rank of tensors: Tao

A tensor $T : [N]^3 \to \mathbb{F}$ has slice rank 1 if there are functions $f : [N] \to \mathbb{F}$ and $g : [N]^2 \to \mathbb{F}$ such that one of the following holds:

$$T(i, j, k) = f(i)g(j, k)$$

Proof sketch of the multicolored sum-free problem

Slice rank of tensors: Tao

A tensor $T : [N]^3 \to \mathbb{F}$ has slice rank 1 if there are functions $f : [N] \to \mathbb{F}$ and $g : [N]^2 \to \mathbb{F}$ such that one of the following holds:

$$T(i,j,k) = f(i)g(j,k)$$

$$T(i,j,k) = f(i,k)g(j)$$

Proof sketch of the multicolored sum-free problem

Slice rank of tensors: Tao

A tensor $T : [N]^3 \to \mathbb{F}$ has slice rank 1 if there are functions $f : [N] \to \mathbb{F}$ and $g : [N]^2 \to \mathbb{F}$ such that one of the following holds:

$$T(i,j,k) = f(i)g(j,k)$$

$$T(i,j,k) = f(i,k)g(j)$$

$$T(i,j,k) = f(i,j)g(k)$$

Proof sketch of the multicolored sum-free problem

Slice rank of tensors: Tao

A tensor $T : [N]^3 \to \mathbb{F}$ has slice rank 1 if there are functions $f : [N] \to \mathbb{F}$ and $g : [N]^2 \to \mathbb{F}$ such that one of the following holds:

$$T(i,j,k) = f(i)g(j,k)$$

$$T(i,j,k) = f(i,k)g(j)$$

$$T(i,j,k) = f(i,j)g(k)$$

Slice rank of general tensor T: minimum number of rank one tensors needed to sum to T.

Proof sketch of the multicolored sum-free problem

Slice rank of tensors: Tao

A tensor $T : [N]^3 \to \mathbb{F}$ has slice rank 1 if there are functions $f : [N] \to \mathbb{F}$ and $g : [N]^2 \to \mathbb{F}$ such that one of the following holds:

$$T(i,j,k) = f(i)g(j,k)$$

$$T(i,j,k) = f(i,k)g(j)$$

$$T(i,j,k) = f(i,j)g(k)$$

Slice rank of general tensor T: minimum number of rank one tensors needed to sum to T.

Claim

Diagonal tensor has rank equal to number of nonzero elements.

Proof sketch of the multicolored sum-free problem

Let M_n^d be the set of monomials of total degree at most d in n variables, and degree less than p in each variable.

Proof sketch of the multicolored sum-free problem

Let M_n^d be the set of monomials of total degree at most d in n variables, and degree less than p in each variable.

Claim

Take $X = \{x^j\}_{j=1}^m$, $Y = \{y^j\}_{j=1}^m$, $Z = \{z^j\}_{j=1}^m$ in $\mathbb{F}_p^n$, as in the multicolored sum-free problem. Then

$$m \leq 3|M_n^{(p-1)n/3}|$$

Proof sketch of the multicolored sum-free problem

Let M_n^d be the set of monomials of total degree at most d in n variables, and degree less than p in each variable.

Claim

Take $X = \{x^j\}_{j=1}^m$, $Y = \{y^j\}_{j=1}^m$, $Z = \{z^j\}_{j=1}^m$ in $\mathbb{F}_p^n$, as in the multicolored sum-free problem. Then

$$m \leq 3|M_n^{(p-1)n/3}|$$

Take a tensor $T : (\mathbb{F}_p^n)^3 \to \mathbb{F}_p$:

$$T(x, y, z) = \prod_{i=1}^{n} (1 - (x_i + y_i + z_i)^{p-1})$$

Proof sketch of the multicolored sum-free problem

Let M_n^d be the set of monomials of total degree at most d in n variables, and degree less than p in each variable.

Claim

Take $X = \{x^j\}_{j=1}^m$, $Y = \{y^j\}_{j=1}^m$, $Z = \{z^j\}_{j=1}^m$ in $\mathbb{F}_p^n$, as in the multicolored sum-free problem. Then

$$m \leq 3|M_n^{(p-1)n/3}|$$

Take a tensor $T : (\mathbb{F}_p^n)^3 \to \mathbb{F}_p$:

$$T(x,y,z) = \prod_{i=1}^{n} (1 - (x_i + y_i + z_i)^{p-1})$$

T is diagonal on $X \times Y \times Z$, so slice rank is at least m, and is at most $3|M_n^{(p-1)n/3}|$.

Szemerédi's Regularity Lemma

Szemerédi's regularity lemma

Roughly speaking, every large graph can be partitioned into a bounded number of roughly equally-sized parts so that the graph is random-like between almost all pairs of parts.

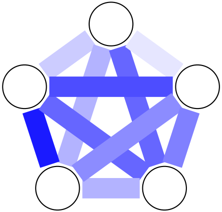

Szemerédi's Regularity Lemma

Szemerédi's regularity lemma

Roughly speaking, every large graph can be partitioned into a bounded number of roughly equally-sized parts so that the graph is random-like between almost all pairs of parts.

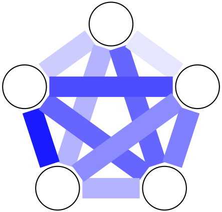

- Rough structural result for all graphs.

Szemerédi's Regularity Lemma

Szemerédi's regularity lemma

Roughly speaking, every large graph can be partitioned into a bounded number of roughly equally-sized parts so that the graph is random-like between almost all pairs of parts.

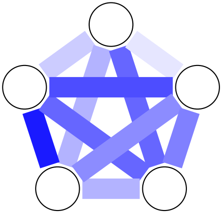

- Rough structural result for all graphs.
- One of the most powerful tools in graph theory.

Szemerédi's Regularity Lemma

Szemerédi's regularity lemma

Roughly speaking, every large graph can be partitioned into a bounded number of roughly equally-sized parts so that the graph is random-like between almost all pairs of parts.

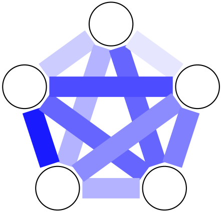

- Rough structural result for all graphs.
- One of the most powerful tools in graph theory.

Triangle Removal Lemma

Triangle Removal Lemma

For every $\varepsilon > 0$ there is $\delta > 0$ such that if a n-vertex graph has at most δn^3 triangles, then we can delete at most εn^2 edges and remove all triangles.

Triangle Removal Lemma

Triangle Removal Lemma

For every $\varepsilon > 0$ there is $\delta > 0$ such that if a n-vertex graph has at most δn^3 triangles, then we can delete at most εn^2 edges and remove all triangles.

Many applications in extremal graph theory, additive number theory, theoretical computer science, and combinatorics.

Triangle Removal Lemma

Triangle Removal Lemma

For every $\varepsilon > 0$ there is $\delta > 0$ such that if a n-vertex graph has at most δn^3 triangles, then we can delete at most εn^2 edges and remove all triangles.

Many applications in extremal graph theory, additive number theory, theoretical computer science, and combinatorics.

Proof uses Szemerédi's regularity lemma, and gives a bound on $1/\delta$ which is a tower of two of height a power of $1/\varepsilon$.

Triangle Removal Lemma

Triangle Removal Lemma

For every $\varepsilon > 0$ there is $\delta > 0$ such that if a n-vertex graph has at most δn^3 triangles, then we can delete at most εn^2 edges and remove all triangles.

Many applications in extremal graph theory, additive number theory, theoretical computer science, and combinatorics.

Proof uses Szemerédi's regularity lemma, and gives a bound on $1/\delta$ which is a tower of two of height a power of $1/\varepsilon$.

Problem (Alon, Erdős, Gowers, Tao)

Find a new proof which gives a better bound.

Triangle Removal Lemma

Triangle Removal Lemma

For every $\varepsilon > 0$ there is $\delta > 0$ such that if a n-vertex graph has at most δn^3 triangles, then we can delete at most εn^2 edges and remove all triangles.

Many applications in extremal graph theory, additive number theory, theoretical computer science, and combinatorics.

Proof uses Szemerédi's regularity lemma, and gives a bound on $1/\delta$ which is a tower of two of height a power of $1/\varepsilon$.

Problem (Alon, Erdős, Gowers, Tao)

Find a new proof which gives a better bound.

Theorem (F.)

We may take $1/\delta$ to be a tower of twos of height $\log 1/\varepsilon$.

Arithmetic Regularity Lemma

Arithmetic Regularity Lemma

Let $A \subset \mathbb{F}_3^n$. The *density of A in S* is $d_A(S) = |A \cap S|/|S|$.

Arithmetic Regularity Lemma

Let $A \subset \mathbb{F}_3^n$. The *density of A in S* is $d_A(S) = |A \cap S|/|S|$.

A translate $S + x \subset \mathbb{F}_3^n$ of a subspace S is ε*-regular* if

$$|d_A(S + x) - d_A(T)| \leq \varepsilon$$

for every codimension 1 affine subspace T of $S + x$.

Arithmetic Regularity Lemma

Let $A \subset \mathbb{F}_3^n$. The *density of A in S* is $d_A(S) = |A \cap S|/|S|$.

A translate $S + x \subset \mathbb{F}_3^n$ of a subspace S is ε*-regular* if

$$|d_A(S + x) - d_A(T)| \leq \varepsilon$$

for every codimension 1 affine subspace T of $S + x$.

A subspace S is ε*-regular* if all but an ε-fraction of the translates of S are ε-regular.

Arithmetic Regularity Lemma

Let $A \subset \mathbb{F}_3^n$. The *density of A in S* is $d_A(S) = |A \cap S|/|S|$.

A translate $S + x \subset \mathbb{F}_3^n$ of a subspace S is ε*-regular* if

$$|d_A(S + x) - d_A(T)| \leq \varepsilon$$

for every codimension 1 affine subspace T of $S + x$.

A subspace S is ε*-regular* if all but an ε-fraction of the translates of S are ε-regular.

Green's arithmetic regularity lemma

For each $\varepsilon > 0$ there is $M(\varepsilon)$ such that for any $A \subset \mathbb{F}_3^n$, there is an ε-regular subspace S of codimension at most $M(\varepsilon)$.

Arithmetic Regularity Lemma

Let $A \subset \mathbb{F}_3^n$. The *density of A in S* is $d_A(S) = |A \cap S|/|S|$.

A translate $S + x \subset \mathbb{F}_3^n$ of a subspace S is ε*-regular* if

$$|d_A(S+x) - d_A(T)| \leq \varepsilon$$

for every codimension 1 affine subspace T of $S + x$.

A subspace S is ε*-regular* if all but an ε-fraction of the translates of S are ε-regular.

Green's arithmetic regularity lemma

For each $\varepsilon > 0$ there is $M(\varepsilon)$ such that for any $A \subset \mathbb{F}_3^n$, there is an ε-regular subspace S of codimension at most $M(\varepsilon)$.

Green, Hosseini-Lovett-Moshkovitz-Shapira:
$M(\varepsilon)$ is a tower of twos of height $\varepsilon^{-O(1)}$.

Arithmetic Triangle Removal Lemma

A *triangle* in $\mathbb{F}_p^n$ is a triple (x, y, z) of points with $x + y + z = 0$.

Arithmetic Triangle Removal Lemma

A *triangle* in $\mathbb{F}_p^n$ is a triple (x, y, z) of points with $x + y + z = 0$.

Green's Arithmetic Triangle Removal Lemma

For every $\varepsilon > 0$ and prime p, there is $\delta > 0$ such that if $X, Y, Z \subset \mathbb{F}_p^n$ with at most δp^{2n} triangles in $X \times Y \times Z$, then we can delete εp^n points and remove all triangles.

Arithmetic Triangle Removal Lemma

A *triangle* in $\mathbb{F}_p^n$ is a triple (x, y, z) of points with $x + y + z = 0$.

Green's Arithmetic Triangle Removal Lemma

For every $\varepsilon > 0$ and prime p, there is $\delta > 0$ such that if $X, Y, Z \subset \mathbb{F}_p^n$ with at most δp^{2n} triangles in $X \times Y \times Z$, then we can delete εp^n points and remove all triangles.

Green's proof uses the arithmetic regularity lemma and gives a bound on $1/\delta$ which is a tower of two of height a power of $1/\varepsilon$.

Arithmetic Triangle Removal Lemma

A *triangle* in $\mathbb{F}_p^n$ is a triple (x, y, z) of points with $x + y + z = 0$.

Green's Arithmetic Triangle Removal Lemma

For every $\varepsilon > 0$ and prime p, there is $\delta > 0$ such that if $X, Y, Z \subset \mathbb{F}_p^n$ with at most δp^{2n} triangles in $X \times Y \times Z$, then we can delete εp^n points and remove all triangles.

Green's proof uses the arithmetic regularity lemma and gives a bound on $1/\delta$ which is a tower of two of height a power of $1/\varepsilon$.

Král'-Serra-Vena: new proof using graph triangle removal lemma.

Arithmetic Triangle Removal Lemma

A *triangle* in $\mathbb{F}_p^n$ is a triple (x, y, z) of points with $x + y + z = 0$.

Green's Arithmetic Triangle Removal Lemma

For every $\varepsilon > 0$ and prime p, there is $\delta > 0$ such that if $X, Y, Z \subset \mathbb{F}_p^n$ with at most δp^{2n} triangles in $X \times Y \times Z$, then we can delete εp^n points and remove all triangles.

Green's proof uses the arithmetic regularity lemma and gives a bound on $1/\delta$ which is a tower of two of height a power of $1/\varepsilon$.

Král'-Serra-Vena: new proof using graph triangle removal lemma.

Problem (Green)

Improve the bound in the arithmetic triangle removal lemma.

Král', Serra, Vena proof

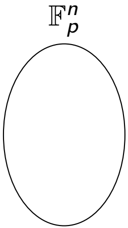

Král', Serra, Vena proof

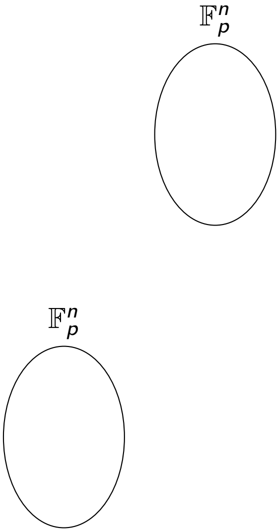

Král', Serra, Vena proof

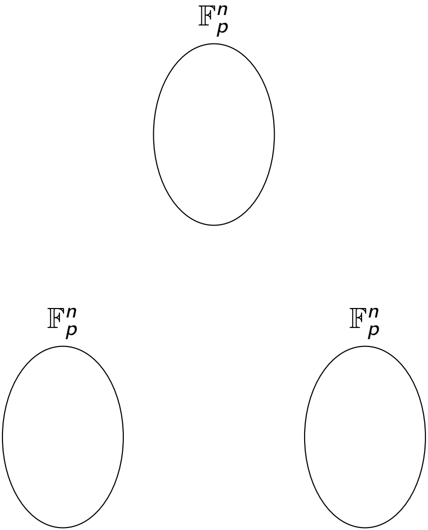

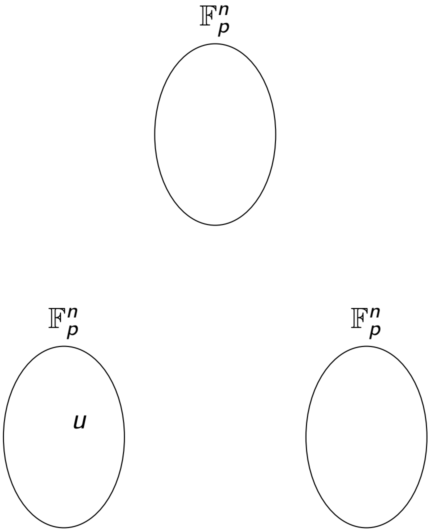
$\mathbb{F}_p^n$
$\mathbb{F}_p^n$
u
$\mathbb{F}_p^n$

Kráľ, Serra, Vena proof

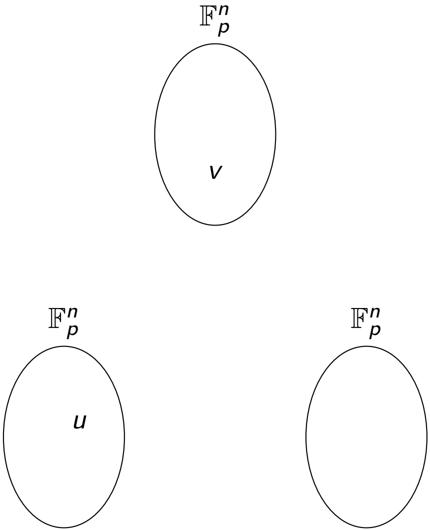

Král', Serra, Vena proof

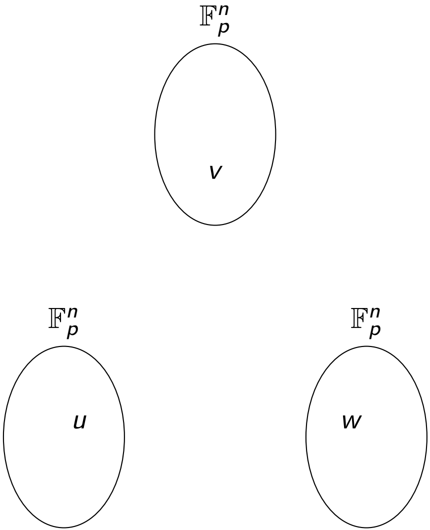

Král', Serra, Vena proof

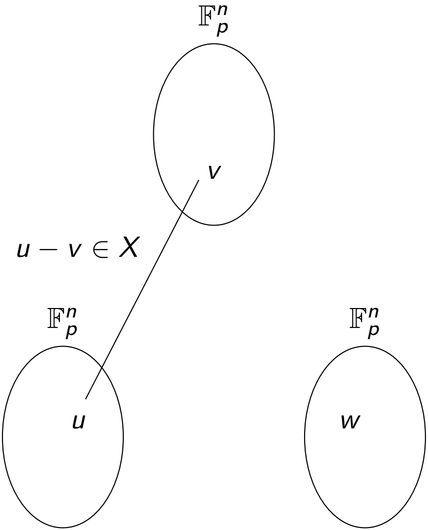

Král', Serra, Vena proof

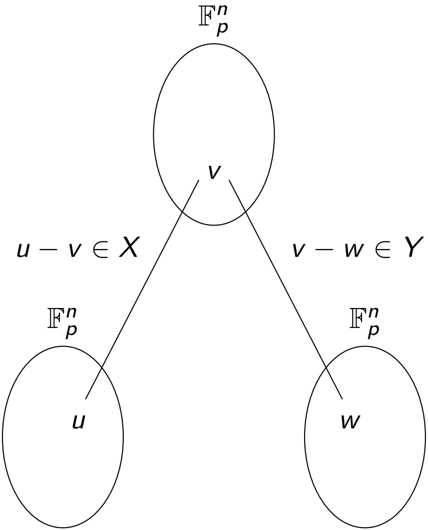

Král', Serra, Vena proof

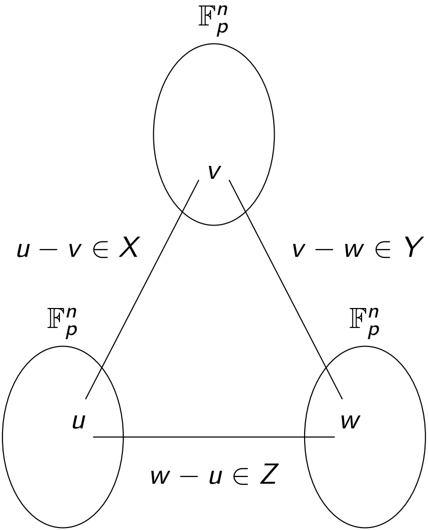

Král', Serra, Vena proof

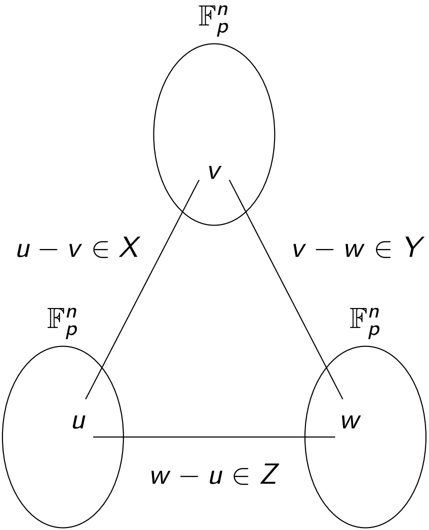

Triangle $x + y + z = 0$ corresponds to $N := p^n$ triangles in the graph, and vice versa.

Král', Serra, Vena proof

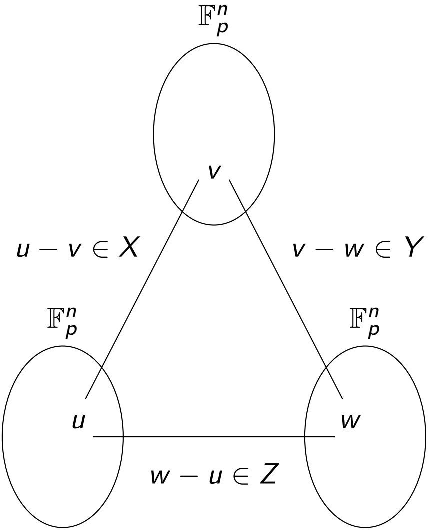

Triangle $x + y + z = 0$ corresponds to $N := p^n$ triangles in the graph, and vice versa.

Thus, there are at most δN^3 triangles.

Král', Serra, Vena proof

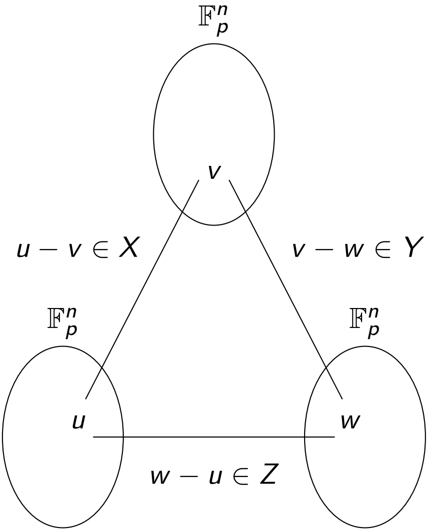

Triangle $x + y + z = 0$ corresponds to $N := p^n$ triangles in the graph, and vice versa.

Thus, there are at most δN^3 triangles.

Can remove εN^2 edges and get rid of all triangles.

Král', Serra, Vena proof

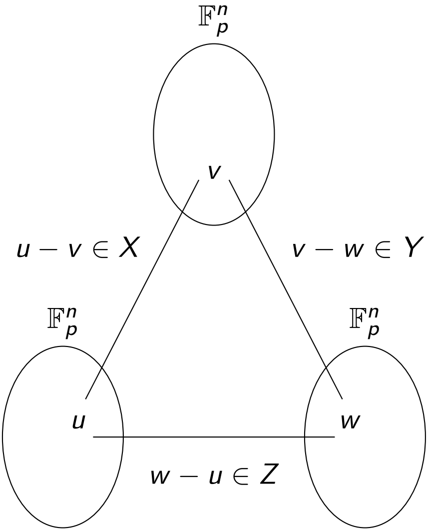

Triangle $x + y + z = 0$ corresponds to $N := p^n$ triangles in the graph, and vice versa.

Thus, there are at most δN^3 triangles.

Can remove εN^2 edges and get rid of all triangles.

Remove x from X, Y, or Z if at least $N/3$ edges corresponding to it are removed.

Arithmetic Triangle Removal Lemma

A *triangle* in $\mathbb{F}_p^n$ is a triple (x, y, z) of points with $x + y + z = 0$.

Green's Arithmetic Triangle Removal Lemma

For every $\varepsilon > 0$ and prime p, there is $\delta > 0$ such that if $X, Y, Z \in \mathbb{F}_p^n$ with at most δp^{2n} triangles in $X \times Y \times Z$, then we can delete εp^n points and remove all triangles.

Arithmetic Triangle Removal Lemma

A *triangle* in $\mathbb{F}_p^n$ is a triple (x, y, z) of points with $x + y + z = 0$.

Green's Arithmetic Triangle Removal Lemma

For every $\varepsilon > 0$ and prime p, there is $\delta > 0$ such that if $X, Y, Z \in \mathbb{F}_p^n$ with at most δp^{2n} triangles in $X \times Y \times Z$, then we can delete εp^n points and remove all triangles.

Much further work on bounds: Hatami-Sachdeva-Tulsiani, Bhattacharyya-Xie, Fu-Kleinberg, Haviv-Xie.

Arithmetic Triangle Removal Lemma

A *triangle* in $\mathbb{F}_p^n$ is a triple (x, y, z) of points with $x + y + z = 0$.

Green's Arithmetic Triangle Removal Lemma

For every $\varepsilon > 0$ and prime p, there is $\delta > 0$ such that if $X, Y, Z \in \mathbb{F}_p^n$ with at most δp^{2n} triangles in $X \times Y \times Z$, then we can delete εp^n points and remove all triangles.

Much further work on bounds: Hatami-Sachdeva-Tulsiani, Bhattacharyya-Xie, Fu-Kleinberg, Haviv-Xie.

Theorem (F.-Lovász)

We can take $\delta = (\varepsilon/3)^{C_p}$, where $C_p = 1 + 1/c_p$ is a computable number. The exponent C_p is sharp.

Arithmetic Triangle Removal Lemma

A *triangle* in $\mathbb{F}_p^n$ is a triple (x, y, z) of points with $x + y + z = 0$.

Green's Arithmetic Triangle Removal Lemma

For every $\varepsilon > 0$ and prime p, there is $\delta > 0$ such that if $X, Y, Z \in \mathbb{F}_p^n$ with at most δp^{2n} triangles in $X \times Y \times Z$, then we can delete εp^n points and remove all triangles.

Much further work on bounds: Hatami-Sachdeva-Tulsiani, Bhattacharyya-Xie, Fu-Kleinberg, Haviv-Xie.

Theorem (F.-Lovász)

We can take $\delta = (\varepsilon/3)^{C_p}$, where $C_p = 1 + 1/c_p$ is a computable number. The exponent C_p is sharp.

In particular, $C_2 = 1 + 1/(5/3 - \log_2 3) \approx 13.239$

Arithmetic Triangle Removal Lemma

A *triangle* in $\mathbb{F}_p^n$ is a triple (x, y, z) of points with $x + y + z = 0$.

Green's Arithmetic Triangle Removal Lemma

For every $\varepsilon > 0$ and prime p, there is $\delta > 0$ such that if $X, Y, Z \in \mathbb{F}_p^n$ with at most δp^{2n} triangles in $X \times Y \times Z$, then we can delete εp^n points and remove all triangles.

Much further work on bounds: Hatami-Sachdeva-Tulsiani, Bhattacharyya-Xie, Fu-Kleinberg, Haviv-Xie.

Theorem (F.-Lovász)

We can take $\delta = (\varepsilon/3)^{C_p}$, where $C_p = 1 + 1/c_p$ is a computable number. The exponent C_p is sharp.

In particular, $C_2 = 1 + 1/(5/3 - \log_2 3) \approx 13.239$ and $C_3 = 1 + 1/c_3$ where $c_3 = 1 - \frac{\log b}{\log 3}$, $b = a^{-2/3} + a^{1/3} + a^{4/3}$, and $a = \frac{\sqrt{33}-1}{8}$, so $C_3 \approx 13.901$.

Arithmetic Triangle Removal Lemma

A *triangle* in $\mathbb{F}_p^n$ is a triple (x, y, z) of points with $x + y + z = 0$.

Green's Arithmetic Triangle Removal Lemma

For every $\varepsilon > 0$ and prime p, there is $\delta > 0$ such that if $X, Y, Z \in \mathbb{F}_p^n$ with at most δp^{2n} triangles in $X \times Y \times Z$, then we can delete εp^n points and remove all triangles.

Much further work on bounds: Hatami-Sachdeva-Tulsiani, Bhattacharyya-Xie, Fu-Kleinberg, Haviv-Xie.

Theorem (F.-Lovász)

We can take $\delta = (\varepsilon/3)^{C_p}$, where $C_p = 1 + 1/c_p$ is a computable number. The exponent C_p is sharp.

In particular, $C_2 = 1 + 1/(5/3 - \log_2 3) \approx 13.239$ and $C_3 = 1 + 1/c_3$ where $c_3 = 1 - \frac{\log b}{\log 3}$, $b = a^{-2/3} + a^{1/3} + a^{4/3}$, and $a = \frac{\sqrt{33}-1}{8}$, so $C_3 \approx 13.901$.

Removal lemma proof sketch

Theorem (F., L. M. Lovász)

With $\delta = (\varepsilon/3)^{C_p}$, if $X, Y, Z \subset \mathbb{F}_p^n$ have at most δp^{2n} triangles in $X \times Y \times Z$, then we can delete εp^n points and remove all triangles.

Removal lemma proof sketch

Theorem (F., L. M. Lovász)

With $\delta = (\varepsilon/3)^{C_p}$, if $X, Y, Z \subset \mathbb{F}_p^n$ have at most δp^{2n} triangles in $X \times Y \times Z$, then we can delete εp^n points and remove all triangles.

Goal 1

With $\delta = \varepsilon^{C_p}$, the union of any εN disjoint triangles with elements red, yellow, blue have $\geq \delta N^2$ rainbow triangles.

Removal lemma proof sketch

Theorem (F., L. M. Lovász)

With $\delta = (\varepsilon/3)^{C_p}$, if $X, Y, Z \subset \mathbb{F}_p^n$ have at most δp^{2n} triangles in $X \times Y \times Z$, then we can delete εp^n points and remove all triangles.

Goal 1

With $\delta = \varepsilon^{C_p}$, the union of any εN disjoint triangles with elements red, yellow, blue have $\geq \delta N^2$ rainbow triangles.

Goal 2

With $\delta = \varepsilon^{C_p + o(1)}$, the union of any εN disjoint triangles with elements red, yellow, blue have $\geq \delta N^2$ rainbow triangles.

Removal lemma proof sketch

Theorem (F., L. M. Lovász)

With $\delta = (\varepsilon/3)^{C_p}$, if $X, Y, Z \subset \mathbb{F}_p^n$ have at most δp^{2n} triangles in $X \times Y \times Z$, then we can delete εp^n points and remove all triangles.

Goal 1

With $\delta = \varepsilon^{C_p}$, the union of any εN disjoint triangles with elements red, yellow, blue have $\geq \delta N^2$ rainbow triangles.

Goal 2

With $\delta = \varepsilon^{C_p+o(1)}$, the union of any εN disjoint triangles with elements red, yellow, blue have $\geq \delta N^2$ rainbow triangles.

Goal 3

With $\delta = \varepsilon^{C_p+o(1)}$, if we have εN disjoint rainbow triangles with each element in $\approx \beta N$ rainbow triangles, then $\beta \geq \delta/\varepsilon$.

Arithmetic triangle removal lemma proof idea

Arithmetic triangle removal lemma proof idea

Goal 3

With $\delta = \varepsilon^{C_p + o(1)}$, if we have εN disjoint rainbow triangles with each element in $\approx \beta N$ rainbow triangles, then $\beta \geq \delta / \varepsilon$.

Arithmetic triangle removal lemma proof idea

Goal 3

With $\delta = \varepsilon^{C_p + o(1)}$, if we have εN disjoint rainbow triangles with each element in $\approx \beta N$ rainbow triangles, then $\beta \geq \delta/\varepsilon$.

Sample a random affine subspace S with $|S| \approx 1/\beta$ elements.

Arithmetic triangle removal lemma proof idea

Goal 3

With $\delta = \varepsilon^{C_p + o(1)}$, if we have εN disjoint rainbow triangles with each element in $\approx \beta N$ rainbow triangles, then $\beta \geq \delta/\varepsilon$.

Sample a random affine subspace S with $|S| \approx 1/\beta$ elements.

A rainbow triangle is *good* if each of its elements are in exactly one rainbow triangle in S.

Arithmetic triangle removal lemma proof idea

Goal 3

With $\delta = \varepsilon^{C_p + o(1)}$, if we have εN disjoint rainbow triangles with each element in $\approx \beta N$ rainbow triangles, then $\beta \geq \delta/\varepsilon$.

Sample a random affine subspace S with $|S| \approx 1/\beta$ elements.

A rainbow triangle is *good* if each of its elements are in exactly one rainbow triangle in S.

With positive probability, the densities of X, Y, Z in S are $\approx \varepsilon$ and a constant fraction of the elements are in good rainbow triangles.

Arithmetic triangle removal lemma proof idea

Goal 3

With $\delta = \varepsilon^{C_p + o(1)}$, if we have εN disjoint rainbow triangles with each element in $\approx \beta N$ rainbow triangles, then $\beta \geq \delta/\varepsilon$.

Sample a random affine subspace S with $|S| \approx 1/\beta$ elements.

A rainbow triangle is *good* if each of its elements are in exactly one rainbow triangle in S.

With positive probability, the densities of X, Y, Z in S are $\approx \varepsilon$ and a constant fraction of the elements are in good rainbow triangles.

From the multicolor sum-free theorem

$$\varepsilon \ll |S|^{-c_p} \approx (1/\beta)^{-c_p},$$

which gives $\delta \leq \varepsilon^{C_p + o(1)}$.

Progressions with popular differences

Progressions with popular differences

Theorem (Green)

$\forall\ \varepsilon > 0$ there is a least $n(\varepsilon)$ such that if $n \geq n(\varepsilon)$, then $\forall\ A \subset \mathbb{F}_3^n$ of density α, there is a nonzero d such that the density of 3-term arithmetic progressions with common difference d is at least $\alpha^3 - \varepsilon$.

Progressions with popular differences

Theorem (Green)

$\forall\ \varepsilon > 0$ there is a least $n(\varepsilon)$ such that if $n \geq n(\varepsilon)$, then $\forall\ A \subset \mathbb{F}_3^n$ of density α, there is a nonzero d such that the density of 3-term arithmetic progressions with common difference d is at least $\alpha^3 - \varepsilon$.

Quiz

How large is $n(\varepsilon)$?

Progressions with popular differences

Theorem (Green)

$\forall\ \varepsilon > 0$ there is a least $n(\varepsilon)$ such that if $n \geq n(\varepsilon)$, then $\forall\ A \subset \mathbb{F}_3^n$ of density α, there is a nonzero d such that the density of 3-term arithmetic progressions with common difference d is at least $\alpha^3 - \varepsilon$.

Quiz

How large is $n(\varepsilon)$?

(a) $\Theta(\log(1/\varepsilon))$

(b) $\varepsilon^{-\Theta(1)}$

(c) $2^{\varepsilon^{-\Theta(1)}}$

(d) $\mathrm{Tower}\left(\Theta\left(\log(1/\varepsilon)\right)\right)$

(e) $\mathrm{Tower}\left(\Theta\left((1/\varepsilon)^{\Theta(1)}\right)\right)$

Progressions with popular differences

Theorem (Green)

$\forall\ \varepsilon > 0$ there is a least $n(\varepsilon)$ such that if $n \geq n(\varepsilon)$, then $\forall\ A \subset \mathbb{F}_3^n$ of density α, there is a nonzero d such that the density of 3-term arithmetic progressions with common difference d is at least $\alpha^3 - \varepsilon$.

Quiz

How large is $n(\varepsilon)$?

(a) $\Theta(\log(1/\varepsilon))$

(b) $\varepsilon^{-\Theta(1)}$

(c) $2^{\varepsilon^{-\Theta(1)}}$

(d) $\mathrm{Tower}\left(\Theta\left(\log(1/\varepsilon)\right)\right)$

(e) $\mathrm{Tower}\left(\Theta\left((1/\varepsilon)^{\Theta(1)}\right)\right)$

Progressions with popular differences

Theorem (Green)

$\forall\ \varepsilon > 0$ there is a least $n(\varepsilon)$ such that if $n \geq n(\varepsilon)$, then $\forall\ A \subset \mathbb{F}_3^n$ of density α, there is a nonzero d such that the density of 3-term arithmetic progressions with common difference d is at least $\alpha^3 - \varepsilon$.

Theorem (F.-Pham)

$$n(\varepsilon) = \mathrm{Tower}\left(\Theta\left(\log(1/\varepsilon)\right)\right)$$

This is the first application of a regularity lemma where a tower-type bound is shown to be needed.

Progressions with popular differences

Theorem (Green)

$\forall\ \varepsilon > 0$ there is a least $n(\varepsilon)$ such that if $n \geq n(\varepsilon)$, then $\forall\ A \subset \mathbb{F}_3^n$ of density α, there is a nonzero d such that the density of 3-term arithmetic progressions with common difference d is at least $\alpha^3 - \varepsilon$.

Theorem (F.-Pham)

$$n(\varepsilon) = \mathrm{Tower}\left(\Theta\left(\log(1/\varepsilon)\right)\right)$$

This is the first application of a regularity lemma where a tower-type bound is shown to be needed.

Theorem* (F.-Pham-Zhao)

A similar result holds in abelian groups and in $[N]$.

Half the random bound

Definition

Let $n'(\alpha)$ be the least integer such that if $n \geq n'(\alpha)$, then for every $A \subset \mathbb{F}_{29}^n$ of density α, there is a nonzero d such that the density of 3-term APs with common difference d is at least $\alpha^3/2$.

Half the random bound

Definition

Let $n'(\alpha)$ be the least integer such that if $n \geq n'(\alpha)$, then for every $A \subset \mathbb{F}_{29}^n$ of density α, there is a nonzero d such that the density of 3-term APs with common difference d is at least $\alpha^3/2$.

Quiz

How large is $n'(\alpha)$?

Half the random bound

Definition

Let $n'(\alpha)$ be the least integer such that if $n \geq n'(\alpha)$, then for every $A \subset \mathbb{F}_{29}^n$ of density α, there is a nonzero d such that the density of 3-term APs with common difference d is at least $\alpha^3/2$.

Quiz

How large is $n'(\alpha)$?

(a) $\Theta(\log(1/\alpha))$

(b) $\alpha^{-\Theta(1)}$

(c) $2^{\alpha^{-\Theta(1)}}$

(d) $\mathrm{Tower}\left(\Theta\left(\log\log(1/\alpha)\right)\right)$

(e) $\mathrm{Tower}\left(\Theta\left(\log(1/\alpha)\right)\right)$

Half the random bound

Definition

Let $n'(\alpha)$ be the least integer such that if $n \geq n'(\alpha)$, then for every $A \subset \mathbb{F}_{29}^n$ of density α, there is a nonzero d such that the density of 3-term APs with common difference d is at least $\alpha^3/2$.

Quiz

How large is $n'(\alpha)$?

(a) $\Theta(\log(1/\alpha))$

(b) $\alpha^{-\Theta(1)}$

(c) $2^{\alpha^{-\Theta(1)}}$

(d) $\mathrm{Tower}\left(\Theta\left(\log\log(1/\alpha)\right)\right)$

(e) $\mathrm{Tower}\left(\Theta\left(\log(1/\alpha)\right)\right)$

Half the random bound

Definition

Let $n'(\alpha)$ be the least integer such that if $n \geq n'(\alpha)$, then for every $A \subset \mathbb{F}_{29}^n$ of density α, there is a nonzero d such that the density of 3-term APs with common difference d is at least $\alpha^3/2$.

Theorem (F.-Pham)

$$n'(\alpha) = \mathrm{Tower}\left(\Theta\left(\log\log(1/\alpha)\right)\right)$$

Multidimensional cap set problem

Inspired by an idea of David Fox on the game SET.

Multidimensional cap set problem

Inspired by an idea of David Fox on the game SET.

Definition

Let $r(n, d)$ be the maximum $|A|$ over all $A \subset \mathbb{F}_3^n$ which contains no d-dimensional afffine subspace.

Multidimensional cap set problem

Inspired by an idea of David Fox on the game SET.

Definition

Let $r(n, d)$ be the maximum $|A|$ over all $A \subset \mathbb{F}_3^n$ which contains no d-dimensional afffine subspace.

Theorem:

For $N = 3^n$, we have $N^{1-(d+1)3^{-d}} \le r(n, d) \le N^{1-13.902^{-d}}$.

Multidimensional cap set problem

Inspired by an idea of David Fox on the game SET.

Definition

Let $r(n,d)$ be the maximum $|A|$ over all $A \subset \mathbb{F}_3^n$ which contains no d-dimensional afffine subspace.

Theorem:

For $N = 3^n$, we have $N^{1-(d+1)3^{-d}} \leq r(n,d) \leq N^{1-13.902^{-d}}$. Hence, the largest dimension of an affine subspace guaranteed in any subset of $\mathbb{F}_3^n$ of density α is $\Theta\left(\log\left(\frac{n}{\log(1/\alpha)}\right)\right)$.

Multidimensional cap set problem

Inspired by an idea of David Fox on the game SET.

Definition

Let $r(n,d)$ be the maximum $|A|$ over all $A \subset \mathbb{F}_3^n$ which contains no d-dimensional afffine subspace.

Theorem:

For $N = 3^n$, we have $N^{1-(d+1)3^{-d}} \leq r(n,d) \leq N^{1-13.902^{-d}}$.
Hence, the largest dimension of an affine subspace guaranteed in any subset of $\mathbb{F}_3^n$ of density α is $\Theta\left(\log\left(\frac{n}{\log(1/\alpha)}\right)\right)$.

Question:

Asymptotics?

Open Problems

Better estimate the bound on the cap set problem.

Open Problems

Better estimate the bound on the cap set problem.

Prove good estimates for the Green-Tao analogue of Green's popular difference theorem for 4-term APs.

Open Problems

Better estimate the bound on the cap set problem.

Prove good estimates for the Green-Tao analogue of Green's popular difference theorem for 4-term APs.

Obtain reasonable bounds on Roth's theorem and the arithmetic triangle removal lemma in other abelian groups.

Open Problems

Better estimate the bound on the cap set problem.

Prove good estimates for the Green-Tao analogue of Green's popular difference theorem for 4-term APs.

Obtain reasonable bounds on Roth's theorem and the arithmetic triangle removal lemma in other abelian groups.

Better estimate the bounds on higher dimensional cap sets.

Open Problems

Better estimate the bound on the cap set problem.

Prove good estimates for the Green-Tao analogue of Green's popular difference theorem for 4-term APs.

Obtain reasonable bounds on Roth's theorem and the arithmetic triangle removal lemma in other abelian groups.

Better estimate the bounds on higher dimensional cap sets.

Extend the new cap set theorem to longer arithmetic progressions.

Thank you!