Applying Quadratic Methods

W. T. Gowers

University of Cambridge

April 10, 2008

Revision

Definition

$$\|f\|_{U^2}^4 = \mathbb{E}_{x,a,b} f(x)\overline{f(x+a)f(x+b)}f(x+a+b)$$

Definition

$$\begin{aligned}\|f\|_{U^3}^8 &= \mathbb{E}_{x,a,b,c} f(x)\overline{f(x+a)f(x+b)}f(x+a+b)\overline{f(x+c)} \\ &\quad f(x+a+c)f(x+b+c)\overline{f(x+a+b+c)}\end{aligned}$$

The expression $\mathbb{E}_{x,d} f_1(x)f_2(x+d)\dots f_k(x+(k-1)d)$ *is robust under small perturbations in the* U^{k-1} *norm.*

How would we like to generalize “linear” Fourier analysis?

A function with large U^2 norm correlates with a linear phase function.

A function with large U^3 norm correlates with a generalized quadratic phase function.

How would we like to generalize "linear" Fourier analysis?

A function with large U^2 norm correlates with a linear phase function.

A function with large U^3 norm correlates with a generalized quadratic phase function.

The trigonometric functions form an orthonormal basis.

The generalized quadratic phase functions do *not* form a basis.

How would we like to generalize "linear" Fourier analysis?

A function with large U^2 norm correlates with a linear phase function.

A function with large U^3 norm correlates with a generalized quadratic phase function.

Every function f with $\|f\|_2 \leq 1$ can be written as a linear combination of a few trigonometric functions plus a small U^2 error.

Can every function f with $\|f\|_2 \leq 1$ can be written as a linear combination of a few generalized quadratic phase functions plus a small U^3 error?

Normed spaces and duality

Let $X = (\mathbb{R}^n, \|.\|)$ be an n-dimensional normed space and let $f \in X$. Then the *dual space* X^* is the space $(\mathbb{R}^n, \|.\|^*)$, where $\|\phi\|^*$ is defined to be $\max\{|\langle f, \phi\rangle| : \|f\| \leq 1\}$.

Trivial lemma: $\langle f, \phi\rangle \leq \|f\|\|\phi\|^*$ *for every* $f, \phi \in \mathbb{R}^n$.

Similar definitions for norms on $\mathbb{C}^n$.

Example: if $\|f\|_p = (\mathbb{E}_x|f(x)|^p)^{1/p}$ then $\|\phi\|_p^* = \|\phi\|_q$, where $\frac{1}{p} + \frac{1}{q} = 1$. (This depends on the normalization used to define the inner product: $\langle f, \phi\rangle = \mathbb{E}_x f(x)\overline{\phi(x)}$.)

The finitary Hahn-Banach theorem

Definition

Let $X = (\mathbb{R}^n, \|.\|)$ be a normed space and let $f \in X$. A *support functional* for f is a non-zero function $\phi \in \mathbb{R}^n$ such that $\langle f, \phi \rangle = \|f\| \|\phi\|^*$.

Theorem

Every function f in a finite-dimensional normed space has a support functional.

"Proof"

WLOG $\|f\| = 1$, so take a tangent plane P at f to the unit ball of $\|.\|$ and define ϕ to be the unique function such that $P = \{g : \langle g, \phi \rangle = 1\}$.

More generally ...

Theorem

Let K be a convex body in $\mathbb{R}^n$ such that $0 \in K$, and let $f \notin K$. Then there is a function $\phi \in \mathbb{R}^n$ such that $\langle f, \phi \rangle \geq 1$ and such that $\langle g, \phi \rangle \leq 1$ for every $g \in K$.

Corollary

Let $K_1, \ldots, K_r$ be closed convex bodies in $\mathbb{R}^n$, each containing 0 and suppose that $f \notin K_1 + \cdots + K_r$. Then there exists a function $\phi \in \mathbb{R}^n$ and non-negative constants $\lambda_1, \ldots, \lambda_r$ such that

- $\lambda_1 + \cdots + \lambda_r = 1$,
- $\langle f, \phi \rangle > 1$
- $\langle g_i, \phi \rangle \leq \lambda_i$ *for every* $i \leq r$ *and every* $g_i \in K_i$.

Proof of the corollary

$f \notin K_1 + \cdots + K_r$. *Want* ϕ *with*
$\langle f, \phi \rangle > 1$; $\lambda_i \geq 0$; $\sum_i \lambda_i \leq 1$; $\langle g_i, \phi \rangle \leq \lambda_i \ \forall g_i \in K_i$.

Proof of the corollary

$f \notin K_1 + \cdots + K_r$. Want ϕ with
$\langle f, \phi\rangle > 1$; $\lambda_i \geq 0$; $\sum_i \lambda_i \leq 1$; $\langle g_i, \phi\rangle \leq \lambda_i \ \forall g_i \in K_i$.

- Easy to see that $K_1 + \cdots + K_r$ is closed and convex.

Proof of the corollary

$f \notin K_1 + \cdots + K_r$. Want ϕ with
$\langle f, \phi \rangle > 1$; $\lambda_i \geq 0$; $\sum_i \lambda_i \leq 1$; $\langle g_i, \phi \rangle \leq \lambda_i \ \forall g_i \in K_i$.

- Easy to see that $K_1 + \cdots + K_r$ is closed and convex.
- Therefore, there is some $\delta > 0$ such that $(1 - \delta)f \notin K_1 + \cdots + K_r$.

Proof of the corollary

$f \notin K_1 + \cdots + K_r$. Want ϕ with
$\langle f, \phi \rangle > 1$; $\lambda_i \geq 0$; $\sum_i \lambda_i \leq 1$; $\langle g_i, \phi \rangle \leq \lambda_i \; \forall g_i \in K_i$.

- Easy to see that $K_1 + \cdots + K_r$ is closed and convex.
- Therefore, there is some $\delta > 0$ such that $(1-\delta)f \notin K_1 + \cdots + K_r$.
- Therefore, there is a ϕ such that $\langle (1-\delta)f, \phi \rangle \geq 1$ and $\langle g, \phi \rangle \leq 1$ for every $g \in K_1 + \cdots + K_r$.

Proof of the corollary

$f \notin K_1 + \cdots + K_r$. Want ϕ with
$\langle f, \phi \rangle > 1$; $\lambda_i \geq 0$; $\sum_i \lambda_i \leq 1$; $\langle g_i, \phi \rangle \leq \lambda_i$ $\forall g_i \in K_i$.

- Easy to see that $K_1 + \cdots + K_r$ is closed and convex.
- Therefore, there is some $\delta > 0$ such that $(1 - \delta)f \notin K_1 + \cdots + K_r$.
- Therefore, there is a ϕ such that $\langle (1 - \delta)f, \phi \rangle \geq 1$ and $\langle g, \phi \rangle \leq 1$ for every $g \in K_1 + \cdots + K_r$.
- Let $\lambda_i = \max\{\langle g_i, \phi \rangle : g_i \in K_i\}$.

Proof of the corollary

$f \notin K_1 + \cdots + K_r$. Want ϕ with
$\langle f, \phi \rangle > 1$; $\lambda_i \geq 0$; $\sum_i \lambda_i \leq 1$; $\langle g_i, \phi \rangle \leq \lambda_i \ \forall g_i \in K_i$.

- Easy to see that $K_1 + \cdots + K_r$ is closed and convex.
- Therefore, there is some $\delta > 0$ such that $(1-\delta)f \notin K_1 + \cdots + K_r$.
- Therefore, there is a ϕ such that $\langle (1-\delta)f, \phi \rangle \geq 1$ and $\langle g, \phi \rangle \leq 1$ for every $g \in K_1 + \cdots + K_r$.
- Let $\lambda_i = \max\{\langle g_i, \phi \rangle : g_i \in K_i\}$.
- Then $\lambda_i \geq 0$ and $\lambda_1 + \cdots + \lambda_r \leq 1$, or we could just pick $g_i \in K_i$ with $\langle g_i, \phi \rangle = \lambda_i$ and we'd have $\langle g_1 + \cdots + g_r, \phi \rangle > 1$.

Simpler version of the corollary

Let $K_1, \ldots, K_r$ be convex bodies in $\mathbb{R}^n$ that all contain 0 and suppose that $f \notin K_1 + \cdots + K_r$. Then there is a function $\phi \in \mathbb{R}^n$ such that $\langle f, \phi \rangle > 1$ and such that $\langle g_i, \phi \rangle \leq 1$ for every $i \leq r$ and every $g_i \in K_i$.

Inverse theorems imply decomposition theorems

Recall the inverse theorem of Green and Tao.

Theorem

Let $\|f\|_\infty \leq 1$ *and* $\|f\|_{U^3} \geq c$*. Then there is a generalized quadratic phase function* Q *of complexity at most* C *such that* $\langle f, Q\rangle \geq c'$*.*

Inverse theorems imply decomposition theorems

Recall the inverse theorem of Green and Tao.

Theorem

Let $\|f\|_\infty \leq 1$ and $\|f\|_{U^3} \geq c$. Then there is a generalized quadratic phase function Q of complexity at most C such that $\langle f, Q\rangle \geq c'$.

We shall deduce from this the following result (from a forthcoming joint paper with Julia Wolf).

Theorem

Let $\|f\|_2 \leq 1$. Then for every $\epsilon > 0$ one can decompose f as a sum

$$\sum_i \lambda_i Q_i + g + h$$

with $\sum_i |\lambda_i| \leq M = M(\epsilon)$, $\|g\|_{U^3} \leq \epsilon$ and $\|h\|_1 \leq \epsilon$. The Q_i are generalized quadratic phase functions of complexity at most $C(\epsilon)$.

Proof of the decomposition theorem

Want $f = \sum_i \lambda_i Q_i + g + h$ *with* $\sum_i |\lambda_i| \leq M$, $\|g\|_{U^3} \leq \epsilon$, $\|h\|_1 \leq \epsilon$.
Given $\|f\|_2 \leq 1$.

Proof of the decomposition theorem

Want $f = \sum_i \lambda_i Q_i + g + h$ *with* $\sum_i |\lambda_i| \leq M$, $\|g\|_{U^3} \leq \epsilon$, $\|h\|_1 \leq \epsilon$.
Given $\|f\|_2 \leq 1$.

- Suppose that no such decomposition exists.

Proof of the decomposition theorem

Want $f = \sum_i \lambda_i Q_i + g + h$ *with* $\sum_i |\lambda_i| \leq M$, $\|g\|_{U^3} \leq \epsilon$, $\|h\|_1 \leq \epsilon$.
Given $\|f\|_2 \leq 1$.

- Suppose that no such decomposition exists.
- Then there is a function $\phi \in \mathbb{C}^n$ such that $\langle f, \phi \rangle > 1$, and such that ϕ is small in the following three respects:

Proof of the decomposition theorem

Want $f = \sum_i \lambda_i Q_i + g + h$ *with* $\sum_i |\lambda_i| \leq M$, $\|g\|_{U^3} \leq \epsilon$, $\|h\|_1 \leq \epsilon$.
Given $\|f\|_2 \leq 1$.

- Suppose that no such decomposition exists.
- Then there is a function $\phi \in \mathbb{C}^n$ such that $\langle f, \phi \rangle > 1$, and such that ϕ is small in the following three respects:
- $\|\phi\|_{U^3}^* \leq \epsilon^{-1}$ (since $\langle g, \phi \rangle \leq 1$ whenever $\|g\|_{U^3} \leq \epsilon$)

Proof of the decomposition theorem

Want $f = \sum_i \lambda_i Q_i + g + h$ *with* $\sum_i |\lambda_i| \leq M$, $\|g\|_{U^3} \leq \epsilon$, $\|h\|_1 \leq \epsilon$.
Given $\|f\|_2 \leq 1$.

- Suppose that no such decomposition exists.
- Then there is a function $\phi \in \mathbb{C}^n$ such that $\langle f, \phi \rangle > 1$, and such that ϕ is small in the following three respects:
- $\|\phi\|_{U^3}^* \leq \epsilon^{-1}$ (since $\langle g, \phi \rangle \leq 1$ whenever $\|g\|_{U^3} \leq \epsilon$)
- $\|\phi\|_\infty \leq \epsilon^{-1}$ (for a very similar reason)

Proof of the decomposition theorem

Want $f = \sum_i \lambda_i Q_i + g + h$ *with* $\sum_i |\lambda_i| \leq M$, $\|g\|_{U^3} \leq \epsilon$, $\|h\|_1 \leq \epsilon$.
Given $\|f\|_2 \leq 1$.

- Suppose that no such decomposition exists.
- Then there is a function $\phi \in \mathbb{C}^n$ such that $\langle f, \phi \rangle > 1$, and such that ϕ is small in the following three respects:
- $\|\phi\|^*_{U^3} \leq \epsilon^{-1}$ (since $\langle g, \phi \rangle \leq 1$ whenever $\|g\|_{U^3} \leq \epsilon$)
- $\|\phi\|_\infty \leq \epsilon^{-1}$ (for a very similar reason)
- $\langle \sum_i \lambda_i Q_i, \phi \rangle \leq 1$ whenever $\sum_i |\lambda_i| \leq M$ and the Q_i are generalized quadratic phase functions of complexity at most $C(\epsilon)$ — which implies that $\langle Q, \phi \rangle \leq M^{-1}$ for every Q.

Proof of the decomposition theorem

Want $f = \sum_i \lambda_i Q_i + g + h$ *with* $\sum_i |\lambda_i| \leq M$, $\|g\|_{U^3} \leq \epsilon$, $\|h\|_1 \leq \epsilon$. *Given* $\|f\|_2 \leq 1$.

- Suppose that no such decomposition exists.
- Then there is a function $\phi \in \mathbb{C}^n$ such that $\langle f, \phi \rangle > 1$, and such that ϕ is small in the following three respects:
- $\|\phi\|_{U^3}^* \leq \epsilon^{-1}$ (since $\langle g, \phi \rangle \leq 1$ whenever $\|g\|_{U^3} \leq \epsilon$)
- $\|\phi\|_\infty \leq \epsilon^{-1}$ (for a very similar reason)
- $\langle \sum_i \lambda_i Q_i, \phi \rangle \leq 1$ whenever $\sum_i |\lambda_i| \leq M$ and the Q_i are generalized quadratic phase functions of complexity at most $C(\epsilon)$ — which implies that $\langle Q, \phi \rangle \leq M^{-1}$ for every Q.
- But $\langle f, \phi \rangle > 1$ and $\|\phi\|_{U^3}^* \leq \epsilon^{-1}$ imply that $\|\phi\|_{U^3} \geq \epsilon$.

Proof of the decomposition theorem

Want $f = \sum_i \lambda_i Q_i + g + h$ *with* $\sum_i |\lambda_i| \leq M$, $\|g\|_{U^3} \leq \epsilon$, $\|h\|_1 \leq \epsilon$. *Given* $\|f\|_2 \leq 1$.

- Suppose that no such decomposition exists.
- Then there is a function $\phi \in \mathbb{C}^n$ such that $\langle f, \phi \rangle > 1$, and such that ϕ is small in the following three respects:
- $\|\phi\|^*_{U^3} \leq \epsilon^{-1}$ (since $\langle g, \phi \rangle \leq 1$ whenever $\|g\|_{U^3} \leq \epsilon$)
- $\|\phi\|_\infty \leq \epsilon^{-1}$ (for a very similar reason)
- $\langle \sum_i \lambda_i Q_i, \phi \rangle \leq 1$ whenever $\sum_i |\lambda_i| \leq M$ and the Q_i are generalized quadratic phase functions of complexity at most $C(\epsilon)$ — which implies that $\langle Q, \phi \rangle \leq M^{-1}$ for every Q.
- But $\langle f, \phi \rangle > 1$ and $\|\phi\|^*_{U^3} \leq \epsilon^{-1}$ imply that $\|\phi\|_{U^3} \geq \epsilon$.
- This contradicts the inverse theorem!

Arithmetic progressions in uniform sets

Definition

A subset $A \subset \mathbb{Z}_N$ of density δ is *uniform of degree* k if $\|A - \delta\mathbf{1}\|_{U^{k+1}}$ is small.

Theorem

If A is uniform of degree $k-1$, then

$$\mathbb{E}_{x,d} A(x)A(x+d)\dots A(x+(k-1)d) \approx \delta^k.$$

Reminder of proof

If we can approximate A by $\delta\mathbf{1}$ in the U^{k-1} norm, then the left-hand side does not change by much if we replace A by $\delta\mathbf{1}$. But if we do that then we get δ^k.

What about more general linear configurations?

Suppose we had an expression such as

$$\mathbb{E}_{x,y,z} A(x-y)A(x+y+z)A(3y-z)A(x+2y-5z)A(z).$$

If A is a random set of density δ then this will be about δ^5. This suggests that there may well be some k such that the above expectation will be about δ^5 if A is uniform of degree k.

What about more general linear configurations?

Suppose we had an expression such as

$$\mathbb{E}_{x,y,z}A(x-y)A(x+y+z)A(3y-z)A(x+2y-5z)A(z).$$

If A is a random set of density δ then this will be about δ^5. This suggests that there may well be some k such that the above expectation will be about δ^5 if A is uniform of degree k.

Green and Tao worked out the most general result that followed from the techniques used to prove the assertion for APs. This analysis led to the notion of the *complexity* of a system of linear forms.

The Green-Tao notion of complexity

Definition

A system of linear forms $L_1, \ldots, L_r$ has *complexity at most k at i* if it is possible to partition the set $\{L_j : j \neq i\}$ into at most $k+1$ subsets such that L_i is not in the linear span of any of those subsets.

The Green-Tao notion of complexity

Definition

A system of linear forms $L_1, \ldots, L_r$ has *complexity at most k at i* if it is possible to partition the set $\{L_j : j \neq i\}$ into at most $k+1$ subsets such that L_i is not in the linear span of any of those subsets.

Example 1. If L_i is the form $x + iy$, $i = 0, 1, \ldots, k-1$, then any two forms span the whole set of forms. Therefore, any partition has to be into singletons. So the complexity of this system is $k-2$ at every individual form.

The Green-Tao notion of complexity

Definition

A system of linear forms $L_1, \ldots, L_r$ has *complexity at most k at i* if it is possible to partition the set $\{L_j : j \neq i\}$ into at most $k+1$ subsets such that L_i is not in the linear span of any of those subsets.

Example 1. If L_i is the form $x + iy$, $i = 0, 1, \ldots, k-1$, then any two forms span the whole set of forms. Therefore, any partition has to be into singletons. So the complexity of this system is $k-2$ at every individual form.

Example 2. For $1 \leq i < j \leq r$ let L_{ij} be the form $x_i + x_j$. If we exclude L_{ij} then we can partition the rest into two sets of forms such that one set never involves x_i and the other never involves x_j. Therefore, the complexity is 1 at every form (a fact implicitly exploited by Balog).

The U^{k+1} norm controls systems of complexity k.

Theorem

If a system of forms $L_1, \ldots, L_r$ has complexity at most k at every i, and if $A \subset \mathbb{Z}_N$ is a set of density δ that is uniform of degree k, then

$$\mathbb{E}_{\mathbf{x}} A(L_1(\mathbf{x}))A(L_2(\mathbf{x}))\ldots A(L_r(\mathbf{x})) \approx \delta^r$$

From this we can recover the earlier result about arithmetic progressions.

The U^{k+1} norm controls systems of complexity k.

Theorem

If a system of forms $L_1, \ldots, L_r$ has complexity at most k at every i, and if $A \subset \mathbb{Z}_N$ is a set of density δ that is uniform of degree k, then

$$\mathbb{E}_{\mathbf{x}} A(L_1(\mathbf{x}))A(L_2(\mathbf{x}))\ldots A(L_r(\mathbf{x})) \approx \delta^r$$

From this we can recover the earlier result about arithmetic progressions.

What about the converse?

Polynomial phase functions again

Recall the example that proves that expressions such as

$$\mathbb{E}_{x,d} f_1(x) f_2(x+d) f_3(x+2d) f_4(x+3d)$$

are not robust under small U^2 perturbations.

Polynomial phase functions again

Recall the example that proves that expressions such as

$$\mathbb{E}_{x,d} f_1(x) f_2(x+d) f_3(x+2d) f_4(x+3d)$$

are not robust under small U^2 perturbations.

We took $f_1(x) = e_N(x^2)$, $f_2(x) = e_N(-3x^2)$, $f_3(x) = e_N(3x^2)$, and $f_4(x) = e_N(-x^2)$.

Polynomial phase functions again

Recall the example that proves that expressions such as

$$\mathbb{E}_{x,d} f_1(x) f_2(x+d) f_3(x+2d) f_4(x+3d)$$

are not robust under small U^2 perturbations.

We took $f_1(x) = e_N(x^2)$, $f_2(x) = e_N(-3x^2)$, $f_3(x) = e_N(3x^2)$, and $f_4(x) = e_N(-x^2)$.

Crucial to this was the fact that

$$x^2 - 3(x+d)^2 + 3(x+2d)^2 - (x+3d)^2 \equiv 0.$$

In particular, the squares of the linear forms x, $x+d$, $x+2d$ and $x+3d$ are linearly dependent.

Three easy facts.

Three easy facts.

- The polynomial phase function $e_N(x^k)$ has small U^k norm but its U^{k+1} norm is 1.

Three easy facts.

- The polynomial phase function $e_N(x^k)$ has small U^k norm but its U^{k+1} norm is 1.

- If the linear forms $L_1, L_2, \ldots, L_r$ have kth powers that are linearly dependent, then we can use functions of the form $e_N(ax^k)$ to prove that expressions of the form

$$\mathbb{E}_{\mathbf{x}} f_1(L_1(\mathbf{x})) f_2(L_2(\mathbf{x})) \ldots f_r(L_r(\mathbf{x}))$$

  are not robust under small U^{k+1} perturbations.

Three easy facts.

- The polynomial phase function $e_N(x^k)$ has small U^k norm but its U^{k+1} norm is 1.
- If the linear forms $L_1, L_2, \ldots, L_r$ have kth powers that are linearly dependent, then we can use functions of the form $e_N(ax^k)$ to prove that expressions of the form

$$\mathbb{E}_{\mathbf{x}} f_1(L_1(\mathbf{x})) f_2(L_2(\mathbf{x})) \ldots f_r(L_r(\mathbf{x}))$$

  are not robust under small U^{k+1} perturbations.
- The system x, y, z, $x+y+z$, $x-y+2z$, $x+y-2z$ is square-independent but has complexity 2.

Which U^k norm is needed for which linear configuration?

For example, for which k is it the case that the expression

$$\mathbb{E}_{x,y,z} f_1(x) f_2(y) f_3(z) f_4(x+y+z) f_5(x-2y+z) f_6(x+y-2z)$$

is robust under small U^k perturbations of the f_i?

Which U^k norm is needed for which linear configuration?

For example, for which k is it the case that the expression

$$\mathbb{E}_{x,y,z} f_1(x) f_2(y) f_3(z) f_4(x+y+z) f_5(x-2y+z) f_6(x+y-2z)$$

is robust under small U^k perturbations of the f_i?

- Since the complexity of the system is 2, it's enough if $k \geq 3$.

Which U^k norm is needed for which linear configuration?

For example, for which k is it the case that the expression

$$\mathbb{E}_{x,y,z} f_1(x) f_2(y) f_3(z) f_4(x+y+z) f_5(x-2y+z) f_6(x+y-2z)$$

is robust under small U^k perturbations of the f_i?

- Since the complexity of the system is 2, it's enough if $k \geq 3$.
- However, it is not possible to prove that $k \geq 3$ is necessary using quadratic (or even generalized quadratic) phase functions.

The following theorem is a joint result with Julia Wolf. It applies to functions defined on the group $\mathbb{F}_p^n$ rather than the group $\mathbb{Z}_N$.

Theorem

Let $\mathcal{L}$ be a collection of linear forms of complexity 2. Then the expression

$$\mathbb{E}_{x_1,\dots,x_m} \prod_{L\in\mathcal{L}} f_L(L(x_1,\dots,x_m))$$

is robust under small U^2 perturbations of the functions f_L if and only if the squares of the linear forms $L \in \mathcal{L}$ are linearly independent.

Very very rough idea of proof

Very very rough idea of proof

- First, use the decomposition theorem to decompose each f_i as a combination of the form $\sum_i \lambda_i Q_i + g + h$, where the Q_i are (slightly more general than) generalized quadratic phase functions, $\sum_i |\lambda_i|$ is not too large, and $\|g\|_{U^3}$ and $\|h\|_1$ are small.

Very very rough idea of proof

- First, use the decomposition theorem to decompose each f_i as a combination of the form $\sum_i \lambda_i Q_i + g + h$, where the Q_i are (slightly more general than) generalized quadratic phase functions, $\sum_i |\lambda_i|$ is not too large, and $\|g\|_{U^3}$ and $\|h\|_1$ are small.
- Next, prove that if $\|f\|_{U^2}$ is very small, then the decomposition can be carried out in such a way that the Q_i have "high rank" (in particular, each $\|Q_i\|_{U^2}$ is small).

Very very rough idea of proof

- First, use the decomposition theorem to decompose each f_i as a combination of the form $\sum_i \lambda_i Q_i + g + h$, where the Q_i are (slightly more general than) generalized quadratic phase functions, $\sum_i |\lambda_i|$ is not too large, and $\|g\|_{U^3}$ and $\|h\|_1$ are small.
- Next, prove that if $\|f\|_{U^2}$ is very small, then the decomposition can be carried out in such a way that the Q_i have "high rank" (in particular, each $\|Q_i\|_{U^2}$ is small).
- Then, using these decompositions, *explicitly calculate* the expectation over the product of linear forms. The error terms make small contributions, and because of the square independence, the purely quadratic terms do too.

Very very rough idea of proof

- First, use the decomposition theorem to decompose each f_i as a combination of the form $\sum_i \lambda_i Q_i + g + h$, where the Q_i are (slightly more general than) generalized quadratic phase functions, $\sum_i |\lambda_i|$ is not too large, and $\|g\|_{U^3}$ and $\|h\|_1$ are small.
- Next, prove that if $\|f\|_{U^2}$ is very small, then the decomposition can be carried out in such a way that the Q_i have "high rank" (in particular, each $\|Q_i\|_{U^2}$ is small).
- Then, using these decompositions, *explicitly calculate* the expectation over the product of linear forms. The error terms make small contributions, and because of the square independence, the purely quadratic terms do too.
- Rough philosophy: constructions of quadratic type don't work, but quadratic Fourier analysis shows that they are essentially the only constructions around.

A different way of proving decomposition results

Another way to deduce decomposition theorems from inverse theorems, due to Green and Tao, but with precedents in combinatorics (Szemerédi's regularity lemma) and ergodic theory (the "Furstenberg tower"), is to use *energy arguments*.

A different way of proving decomposition results

Another way to deduce decomposition theorems from inverse theorems, due to Green and Tao, but with precedents in combinatorics (Szemerédi's regularity lemma) and ergodic theory (the "Furstenberg tower"), is to use *energy arguments*.

Rough idea: try to find a partition $\mathcal{B}$ of the group (that is, $\mathbb{Z}_N$ or $\mathbb{F}_p^n$) into sets of "quadratic type," in such a way that $f - \mathbb{E}(f|\mathcal{B})$ is small.

A different way of proving decomposition results

Another way to deduce decomposition theorems from inverse theorems, due to Green and Tao, but with precedents in combinatorics (Szemerédi's regularity lemma) and ergodic theory (the "Furstenberg tower"), is to use *energy arguments*.

Rough idea: try to find a partition $\mathcal{B}$ of the group (that is, $\mathbb{Z}_N$ or $\mathbb{F}_p^n$) into sets of "quadratic type," in such a way that $f - \mathbb{E}(f|\mathcal{B})$ is small.

(Definition: $\mathbb{E}(f|\mathcal{B})(x)$ is the average of f over the cell of $\mathcal{B}$ that contains x. So $\mathbb{E}(f|\mathcal{B})$ is the orthogonal projection of f to the space of functions that are constant on the cells of $\mathcal{B}$.)

Energy increments

Rough idea: try to find a partition $\mathcal{B}$ of the group (that is, $\mathbb{Z}_N$ or $\mathbb{F}_p^n$) into sets of "quadratic type," in such a way that $f - \mathbb{E}(f|\mathcal{B})$ is small.

Energy increments

Rough idea: try to find a partition $\mathcal{B}$ of the group (that is, $\mathbb{Z}_N$ or $\mathbb{F}_p^n$) into sets of "quadratic type," in such a way that $f - \mathbb{E}(f|\mathcal{B})$ is small.

Suppose that $\|f - \mathbb{E}(f|\mathcal{B})\|_{U^3}$ is not small. Since f is bounded and $\mathbb{E}(.|\mathcal{B})$ is an averaging projection, $f - \mathbb{E}(f|\mathcal{B})$ is bounded as well. Therefore, the inverse theorem implies that it correlates with a generalized quadratic phase function Q.

Energy increments

Rough idea: try to find a partition $\mathcal{B}$ of the group (that is, $\mathbb{Z}_N$ or $\mathbb{F}_p^n$) into sets of "quadratic type," in such a way that $f - \mathbb{E}(f|\mathcal{B})$ is small.

Suppose that $\|f - \mathbb{E}(f|\mathcal{B})\|_{U^3}$ is not small. Since f is bounded and $\mathbb{E}(.|\mathcal{B})$ is an averaging projection, $f - \mathbb{E}(f|\mathcal{B})$ is bounded as well. Therefore, the inverse theorem implies that it correlates with a generalized quadratic phase function Q.

Now pass to a refinement $\mathcal{B}'$ of $\mathcal{B}$ on the cells of which Q is (approximately) constant. Then the correlation implies that $\|\mathbb{E}(f|\mathcal{B}')\|_2$ is substantially bigger than $\|\mathbb{E}(f|\mathcal{B})\|_2$.

Energy increments

Rough idea: try to find a partition $\mathcal{B}$ of the group (that is, $\mathbb{Z}_N$ or $\mathbb{F}_p^n$) into sets of "quadratic type," in such a way that $f - \mathbb{E}(f|\mathcal{B})$ is small.

Suppose that $\|f - \mathbb{E}(f|\mathcal{B})\|_{U^3}$ is not small. Since f is bounded and $\mathbb{E}(.|\mathcal{B})$ is an averaging projection, $f - \mathbb{E}(f|\mathcal{B})$ is bounded as well. Therefore, the inverse theorem implies that it correlates with a generalized quadratic phase function Q.

Now pass to a refinement $\mathcal{B}'$ of $\mathcal{B}$ on the cells of which Q is (approximately) constant. Then the correlation implies that $\|\mathbb{E}(f|\mathcal{B}')\|_2$ is substantially bigger than $\|\mathbb{E}(f|\mathcal{B})\|_2$.

The cells of these partitions are "quadratic" in the sense that they are (approximate) level sets of generalized quadratic phase functions.

Comparison of the two methods

Comparison of the two methods

- The energy-increment argument preserves boundedness because it is based on averaging projections, whereas the decomposition using the Hahn-Banach theorem is not necessarily into bounded parts.

Comparison of the two methods

- The energy-increment argument preserves boundedness because it is based on averaging projections, whereas the decomposition using the Hahn-Banach theorem is not necessarily into bounded parts.
- In both methods, technical difficulties arise if the quadratic sets/functions have low rank. It seems that dealing with these difficulties is more expensive for the energy-increment method.

Comparison of the two methods

- The energy-increment argument preserves boundedness because it is based on averaging projections, whereas the decomposition using the Hahn-Banach theorem is not necessarily into bounded parts.
- In both methods, technical difficulties arise if the quadratic sets/functions have low rank. It seems that dealing with these difficulties is more expensive for the energy-increment method.
- For our purposes, the advantage of Hahn-Banach outweighed the disadvantage. (That is, we obtained a weaker result with a much better bound and were able to apply it.)

Comparison of the two methods

- The energy-increment argument preserves boundedness because it is based on averaging projections, whereas the decomposition using the Hahn-Banach theorem is not necessarily into bounded parts.
- In both methods, technical difficulties arise if the quadratic sets/functions have low rank. It seems that dealing with these difficulties is more expensive for the energy-increment method.
- For our purposes, the advantage of Hahn-Banach outweighed the disadvantage. (That is, we obtained a weaker result with a much better bound and were able to apply it.)
- In one respect, our result was stronger: our initial assumption was that $\|f\|_2 \leq 1$. But we don't use this and we pay for it with an L_1 error term.

Comparison of the two methods

- The energy-increment argument preserves boundedness because it is based on averaging projections, whereas the decomposition using the Hahn-Banach theorem is not necessarily into bounded parts.
- In both methods, technical difficulties arise if the quadratic sets/functions have low rank. It seems that dealing with these difficulties is more expensive for the energy-increment method.
- For our purposes, the advantage of Hahn-Banach outweighed the disadvantage. (That is, we obtained a weaker result with a much better bound and were able to apply it.)
- In one respect, our result was stronger: our initial assumption was that $\|f\|_2 \leq 1$. But we don't use this and we pay for it with an L_1 error term.
- The Hahn-Banach method is often simpler and more direct.

The main difficulty

Recall the decomposition theorem:

If f is any bounded function and $\epsilon > 0$ then we can write it as $\sum_i \lambda_i Q_i + g + h$ such that $\sum_i |\lambda_i| \leq M(\epsilon)$, each Q_i is a generalized quadratic phase function, $\|g\|_{U^3} \leq \epsilon$ and $\|h\|_1 \leq \epsilon$.

The main difficulty

Recall the decomposition theorem:

If f is any bounded function and $\epsilon > 0$ then we can write it as $\sum_i \lambda_i Q_i + g + h$ such that $\sum_i |\lambda_i| \leq M(\epsilon)$, each Q_i is a generalized quadratic phase function, $\|g\|_{U^3} \leq \epsilon$ and $\|h\|_1 \leq \epsilon$.

We needed to know something more: that if $\|f\|_{U^2}$ is small enough, then the functions Q_i have large rank, which is equivalent to saying that they too have small U^2 norms.

The main difficulty

Recall the decomposition theorem:

If f is any bounded function and $\epsilon > 0$ then we can write it as $\sum_i \lambda_i Q_i + g + h$ such that $\sum_i |\lambda_i| \leq M(\epsilon)$, each Q_i is a generalized quadratic phase function, $\|g\|_{U^3} \leq \epsilon$ and $\|h\|_1 \leq \epsilon$.

We needed to know something more: that if $\|f\|_{U^2}$ is small enough, then the functions Q_i have large rank, which is equivalent to saying that they too have small U^2 norms.

It can be shown that if Q has small rank, then $\|Q\|_{U^2}^*$ is small, so $\langle f, Q \rangle$ is small. Therefore, it ought not to be necessary to use Q in a decomposition of f. However, proving this is surprisingly tricky.

Uniform functions need just high-rank quadratics

The rough idea is this. First, write $f = \sum_i \lambda_i Q_i + g + h$ as before. Then find an interval $I = [r, s(r)]$ such that r is fairly large and not many Q_i have rank in I.

Uniform functions need just high-rank quadratics

The rough idea is this. First, write $f = \sum_i \lambda_i Q_i + g + h$ as before. Then find an interval $I = [r, s(r)]$ such that r is fairly large and not many Q_i have rank in I.

Then write $f = f_L + f_H + g + h$. Here, f_L is "low-rank part" and f_H is "high-rank part" of the quadratic decomposition. (The intermediate part is absorbed into the error term.)

Uniform functions need just high-rank quadratics

The rough idea is this. First, write $f = \sum_i \lambda_i Q_i + g + h$ as before. Then find an interval $I = [r, s(r)]$ such that r is fairly large and not many Q_i have rank in I.

Then write $f = f_L + f_H + g + h$. Here, f_L is "low-rank part" and f_H is "high-rank part" of the quadratic decomposition. (The intermediate part is absorbed into the error term.)

Then $\|f_L\|_{U^2}^*$ is not too large. Standard techniques allow us to convolve with a function β in such a way that $\beta * f_L \approx f_L$ in L_2. But that $\|\beta * f\|_2$ and $\|\beta * f_H\|_2$ are small (since $\|f\|_{U^2}$ and $\|f_H\|_{U^2}$ are small). Also, $\|\beta * g\|_{U^3} \leq \|g\|_{U^3}$ and $\|\beta * h\|_1 \leq \|h\|_1$.

Uniform functions need just high-rank quadratics

The rough idea is this. First, write $f = \sum_i \lambda_i Q_i + g + h$ as before. Then find an interval $I = [r, s(r)]$ such that r is fairly large and not many Q_i have rank in I.

Then write $f = f_L + f_H + g + h$. Here, f_L is "low-rank part" and f_H is "high-rank part" of the quadratic decomposition. (The intermediate part is absorbed into the error term.)

Then $\|f_L\|_{U^2}^*$ is not too large. Standard techniques allow us to convolve with a function β in such a way that $\beta * f_L \approx f_L$ in L_2. But that $\|\beta * f\|_2$ and $\|\beta * f_H\|_2$ are small (since $\|f\|_{U^2}$ and $\|f_H\|_{U^2}$ are small). Also, $\|\beta * g\|_{U^3} \leq \|g\|_{U^3}$ and $\|\beta * h\|_1 \leq \|h\|_1$.

It follows that f_L can be absorbed into the error terms.

A clustering argument

The following is an oversimplification of a true statement:

If Q and Q' are generalized quadratic phase functions, then either $\langle Q, Q' \rangle$ is small or $\|Q\overline{Q'}\|_{U^2}^*$ is not too large.

In the second case, we think of Q and Q' as "linearly related".

A clustering argument

The following is an oversimplification of a true statement:

If Q and Q' are generalized quadratic phase functions, then either $\langle Q, Q' \rangle$ is small or $\|Q\overline{Q'}\|_{U^2}^*$ is not too large.

In the second case, we think of Q and Q' as "linearly related".

We also prove a lemma along the following lines.

A sum of the form $\sum_i \lambda_i Q_i$ such that $\sum_i |\lambda_i| \leq M$ can be split into "clusters" of linearly related Q_i plus a remainder that is small in L_2.

This result is basically a simple statement about vectors in Hilbert spaces combined with the fact that if Q and Q' are not approximately orthogonal then they are linearly related.

A more precise decomposition theorem

These ideas lead to the following refinement of the decomposition theorem.

Theorem

Let f be a function with $\|f\|_2 \leq 1$ and let $\epsilon > 0$. Then it is possible to write f as a sum of the form $\sum_{i=1}^{k} Q_i U_i + g + h$, such that $k \leq k(\epsilon)$, $\sum_i \|U_i\|_{U^2}^* \leq M(\epsilon)$, $\|g\|_{U^3} \leq \epsilon$ and $\|h\|_1 \leq \epsilon$.