LECTURE NOTES 10 FOR 247B

TERENCE TAO

1. Representation theory for finite non-abelian groups

In last week's notes we obtained a satisfactory theory of the Fourier transform on finite abelian groups, and then more generally for locally compact abelian groups. Now let's do it for finite non-abelian groups. (The case of infinite non-abelian groups is significantly more complicated, and will only be discussed in an *ad hoc* manner here.) Thus we let $G = (G, \cdot)$ be a finite multiplicative group, which we again give the discrete topology and σ-algebra, and normalised counting measure

$$\int_G f(x)\ dx := \frac{1}{\#G} \sum_{x \in G} f(x).$$

Thus we have the Hilbert space $L^2(G)$ as before.

One might naively hope that the theory of multiplicative characters $\chi : G \to S^1$, which worked so well in the abelian case, carries over easily to the non-abelian case. For instance, it is not hard to adapt the abelian theory to show that any two multiplicative characters are orthogonal, which is an encouraging start. However, a basic problem arises: there are just not enough multiplicative characters to go around. Indeed, if $\chi : G \to S^1$ is a multiplicative character and $x, y \in G$, and $[x, y] := xyx^{-1}y^{-1}$ is the commutator of x and y, then

$$\chi([x, y]) = \chi(x)\chi(y)\chi(x)^{-1}\chi(y)^{-1} = 1.$$

Thus if $[G, G]$ is the group generated by all the commutators $[x, y]$ (which is easily verified to be normal), then every multiplicative character annihilates $[G, G]$, and thus descends to a character on the group $G/[G, G]$ (which is easily verified to be abelian). Using the abelian theory, we thus see that the number of multiplicative characters is only $\#G/\#[G, G]$, which is not enough to span $L^2(G)$ in the non-abelian case. Something more is needed. (For instance, if $G = S_n$, then $[G, G] = A_n$, and so there are only $\#S_n/\#A_n = 2$ characters - the trivial character 1 and the signature $\mathrm{sgn} : S_n \to \{-1, +1\}$.)

To obtain a Fourier theory on G we will use the representation theoretic approach (which seems to be the best approach we currently have in the noncommutative setting).

Definition 1.1 (Unitary representations)**.** Let G be a finite multiplicative group. A (finite dimensional) *unitary representation* is a finite-dimensional[1] complex Hilbert space V, together with a homomorphism $\rho : G \to U(V)$ from G to the unitary group of V, thus for each $x \in G$, $\rho(x) : V \to V$ is unitary, and $\rho(xy) = \rho(x)\rho(y)$ for $x, y \in G$. We define the *dimension* of ρ to be the (complex) dimension of V. (More generally, one often sees V used to denote the representation instead of ρ, although strictly speaking both are abuses of notation, it is the *pair* (ρ, V) which is the representation of G.)

(It turns out that all representations of *finite* groups can be made unitary by adjoining an appropriate Hilbert space structure; see Q1.)

We now give some basic examples of representations. There is the *trivial representation* id_V, in which V is an arbitrary Hilbert space (in particular, it could be $\mathbf{C}$) and $\mathrm{id}_V(x) = \mathrm{id}_V$ is the identity on V for every $x \in G$. More generally, given a multiplicative character χ and an arbitrary Hilbert space V, we have the representation $\chi \mathrm{id}_V$ which assigns the constant multiple $\chi(x)\mathrm{id}_V$ of the identity to each $x \in G$. There is the *regular representation* Trans, in which $V = L^2(G)$ and $\mathrm{Trans}(x)$ is the left translation by x:

$$\mathrm{Trans}(x)f(y) := f(x^{-1}y).$$

(Note that it is only *left*-translation which is a representation: why?) Finally, we note the *zero representation* 0, in which $V = \{0\}$, and $0(x) = 0$. (The zero operator is usually not unitary - *except* when V is itself zero!)

In the previous notes we focused primarily on the regular representation and its subrepresentations, but we shall shortly see that in fact once we understand these special representations, we in fact can understand all other representations relatively easily.

There are two basic binary operations on representations (analogous to natural number addition and multiplication, which corresponds here to the special case where G is trivial). Firstly, Given two representations $\rho_1 : G \to U(V_1)$ and $\rho_2 : G \to U(V_2)$, we can form their *direct sum* $\rho_1 \oplus \rho_2 : G \to U(V_1 \oplus V_2)$ by $\rho_1 \oplus \rho_2(x) := \rho_1(x) \oplus \rho_2(x)$, where $V_1 \oplus V_2$ is the direct sum (i.e. orthogonal sum) of V_1 and V_2, and $\rho_1(x) \oplus \rho_2(x)$ is the block-diagonal operator $\rho_1(x) \oplus \rho_2(x)(v_1, v_2) := (\rho_1(x)v_1, \rho_2(x)v_2)$. We can also form the *tensor product* $\rho_1 \otimes \rho_2 : G \to U(V_1 \otimes V_2)$, where $V_1 \otimes V_2$ is the tensor product of the Hilbert spaces V_1, V_2 (with an orthonormal basis given by $e_{1,i} \otimes e_{2,j}$, where $e_{1,i}$ and $e_{2,j}$ range independently over orthonormal bases of V_1, V_2 respectively), and $\rho_1 \otimes \rho_2(x) := \rho_1(x) \otimes \rho_2(x)$ is the tensor product map, thus

$$(\rho_1(x) \otimes \rho_2(x))(v_1 \otimes v_2) = (\rho_1(x)v_1) \otimes (\rho_2(x)v_2)$$

for any $v_1 \in V_1, v_2 \in V_2$. Thus for instance the zero representation is (up to isomorphism) an identity element for direct sum, while the trivial representation on $\mathbf{C}$ is (again up to isomorphism) the identity element for tensor product. The tensor

[1]One can easily use Zorn's lemma to express infinite-dimensional representations of finite groups as the direct sum of finite-dimensional ones, but we will only need the finite-dimensional theory here.

product of a character representation $\chi_1 \mathrm{id}_{V_1}$ with another $\chi_2 \mathrm{id}_{V_2}$ is $(\chi_1\chi_2)\mathrm{id}_{V_1 \otimes V_2}$. The exact relationship between tensor product and direct sum is a very interesting question - the *Clebsch-Gordan problem* - but we will not dwell on it here.

Problem 1.2. Verify that direct sum and tensor product are commutative and associative up to isomorphism, and that tensor product distributes over direct sum up to isomorphism; this further reinforces the analogy between these operations and addition and multiplication.

A representation ρ is *faithful* when the map $\rho : G \to U(V)$ is injective. Thus for instance the trivial representation is only faithful when G is trivial; at the other extreme, the regular representation is always faithful. Informally, a faithful representation is "big" enough to capture all the behaviour of G, otherwise it is only really capturing the action of a quotient of G. A representation ρ is *free* if there are no fixed points other than the origin; more precisely, for every non-zero $v \in V$, there exists $x \in G$ such that $\rho(x)v \neq v$. Observe that the regular representation is also free. Informally, a free representation contains no trivial component.

Problem 1.3. Let $\rho : G \to U(V)$. Show that there exists a unique normal subgroup H of G and faithful representation $\tilde{\rho} : G/H \to U(V)$ such that ρ is the composition of the quotient map from G to G/H and $\tilde{\rho}$. Also, show that ρ can be uniquely expressed as the direct sum of a free representation and a trivial representation. Then combine these two reductions, and express an arbitrary representation as the direct sum of a trivial representation, and a faithful free representation of a quotient group.

Problem 1.4. Give examples of a representation which is faithful but not free, and vice versa.

A *morphism* from one representation $\rho_1 : G \to U(V_1)$ to another $\rho_2 : G \to U(V_2)$ is a linear map $\phi : V_1 \to V_2$ which intertwines the two representations, thus $\phi \circ \rho_1(x) = \rho_2(x) \circ \phi$. If ϕ is invertible, then the inverse is also a morphism, and ϕ is then said to be an *isomorphism*, and ρ_1 and ρ_2 are *isomorphic*. This is clearly an equivalence relation, and so it is meaningful to talk about a representation ρ obeying some property "up to isomorphism".

A *subrepresentation* of a representation $\rho : G \to U(V)$ is a representation $\rho_W : G \to U(W)$, where W is an *invariant subspace* of V (i.e. W is preserved by all the transformations $\rho(x)$ for $x \in G$), and $\rho|_W(x) : W \to W$ is the restriction of $\rho(x)$ to W (which is automatically unitary). From last week we already saw the importance of invariant subspaces in abstract Fourier analysis. We say that a representation is *irreducible* if is non-zero, and it does not contain a proper non-zero subrepresentation. From the unitary we observe that if W is an invariant subspace in V, then so is the orthogonal complement $W^\perp$; as a consequence we see that irreducibility for representations of finite groups is the same concept as *indecomposability*. Arguing as in the previous week's notes we see that every (finite-dimensional) representation can be expressed as the direct sum of irreducible representations (i.e. all representations of finite groups are *semisimple*). This representation is not unique; consider the identity representation on a Hilbert space V, which has one such decomposition for each orthonormal basis of V.

Problem 1.5. Show that a representation $G \to U(V)$ is irreducible if and only if, for every non-zero $v \in V$, the orbit $\{\rho(x)v : x \in G\}$ of v spans V. Conclude that except for the trivial representation on a one-dimensional vector space, every irreducible representation is free.

The following fundamental lemma allows us to analyse representations in terms of their irreducible components:

Lemma 1.6 (Schur's lemma)**.** *Let $\phi : V \to W$ be a morphism from one representation to another.*

- *If V is irreducible, then ϕ is either injective or zero.*
- *If W is irreducible, then ϕ is either surjective or zero.*
- *If V and W are both irreducible, then ϕ is either an isomorphism or zero.*
- *If $V = W$ is irreducible, then ϕ is a multiple of the identity.*

Proof The range of ϕ is a subrepresentation of W, and is thus either W or $\{0\}$ if W is irreducible. Similarly for the kernel of ϕ. This gives the first two claims, which then easily gives the third. For the fourth, observe from the third claim that ϕ minus any multiple of the identity is either invertible or zero, i.e. the spectrum is a single point, and so ϕ is a multiple of the identity. ■

As a sample application we have

Corollary 1.7 (Irreducible representations are prime)**.** *Suppose that U is an irreducible subrepresentation of $V \oplus W$. Then U is also isomorphic to a subrepresentation of either V or W.*

Proof By Schur's lemma, the projection of U to V is either isomorphic to V or zero, and similarly to W. But the projections cannot both be zero, and the claim follows. ■

By repeating the proof of the fundamental theorem of arithmetic, we conclude a unique factorisation property for representations of finite groups:

Corollary 1.8 (Jordan-Hölder theorem)**.** *Every representation splits as the direct sum of irreducibles, and the isomorphism class of these factors is unique up to permutations. (In particular, there is a well-defined concept of the* multiplicity *of any given irreducible representation V_ξ inside another representation V.)*

Another sample application:

Corollary 1.9. *Let V be an irreducible representation. Then V is isomorphic to a subrepresentation of $L^2(G)$.*

Proof Let $v \in V$ be non-zero, and define a morphism $\phi_{\rho,v} : L^2(G) \to V$ by

$$\phi_{\rho,v}(f) := \int_G f(y)\rho(y)v \, dy;$$

one easily checks that this is indeed a morphism. We split $L^2(G)$ as the direct sum of irreducible representations. By Schur's lemma, each of them is either mapped isomorphically to V by $\phi_{\rho,v}$ or mapped to zero. Since $v \neq 0$, the latter case cannot occur all the time, and the claim follows. ∎

Thus, for instance, we now see from the previous week's notes that the irreducible representations of an abelian group are necessarily one-dimensional, and as a corollary that every representation of a funite abelian group factors as the direct sum of character representations $\chi 1_V$. For instance, the space of functions on $\mathbf{R}$, with the $\mathbf{Z}_2$ action $f(x) \mapsto f(-x)$, splits as the direct sum of the even functions (on which $\mathbf{Z}_2$ acts like the identity character) and odd functions (on which $\mathbf{Z}_2$ acts like the non-identity character).

We now also know that every finite abelian group has only finitely many irreducible representations, up to isomorphism.

One final application of Schur's lemma:

Corollary 1.10 (Ergodic theorem)**.** *Let $\rho_V : G \to U(V)$ and $\rho_W : G \to U(W)$ be irreducible representations, and let $T : V \to W$ be a linear map. Let $\langle T \rangle : V \to W$ be the averaged map*

$$\langle T \rangle := \int_G \rho_W(x) T \rho_V(x)^{-1}\ dx.$$

- *If V and W are not isomorphic, then $\langle T \rangle = 0$.*
- *If $V = W$ and $\rho_V = \rho_W$, then $\langle T \rangle = \frac{\mathrm{tr}_V(T)}{\dim(V)} \mathrm{id}_V$, where $\mathrm{tr}_V(T)$ is the trace of T, defined as*

$$\mathrm{tr}_V(T) := \sum_{j=1}^{\dim(V)} \langle T e_j, e_j \rangle,$$

where e_j ranges over any orthonormal basis of V; one easily verifies that this definition is independent of the choice of such a basis.

Proof One easily verifies that $\langle T \rangle$ is a morphism, so that the first claim follows immediately from Schur's lemma. In the second case, Schur's lemma shows that $\langle T \rangle$ is a multiple of the identity; but then one easily checks that $\langle T \rangle$ has the same trace as T, and the claim follows. ∎

2. The Fourier transform for finite non-abelian groups

Now that we have the basic machinery of representation theory, we can define the Fourier transform properly.

First we need the correct notion of the dual group $\hat{G}$. In the abelian case, this was identifiable with the space of multiplicative characters, or equivalently with the representations of G on $\mathbf{C}$, which up to isomorphism are also the same thing as

the irreducible representations. For non-abelian groups, there are not enough multiplicative characters to go around, but it turns out the irreducible representations are enough.

Thus, let $\hat{G}$ be a set indexing the irreducible representations of G up to isomorphism, thus for each $\xi \in \hat{G}$ we have an irreducible representation $\rho_\xi : G \to U(V_\xi)$, and every irreducible representation is isomorphic to exactly one ρ_ξ. This space does not have an obvious group structure (though we will return to this point a bit later), but is finite thanks to Corollaries 1.9, 1.8.

Given any $\xi \in \hat{G}$ and $f \in L^2(G)$, we define the Fourier coefficient $\hat{f}(\xi) : V_\xi \to V_\xi$ to be the linear transformation

$$\hat{f}(\xi) := \int_G f(x)\rho_\xi(x)\ dx.$$

Note how the Fourier coefficient is now a transformation rather than a complex number. To give some indication as to the terminology "Fourier coefficient" observe that

$$\widehat{f * g}(\xi) = \hat{f}(\xi)\hat{g}(\xi).$$

Note also the non-commutativity on both sides, which shows that one cannot hope for a complex-valued Fourier transform to model convolution a non-abelian group G.

For reasons which will become clearer later, we let $d\xi$ be counting measure on $\hat{G}$ weighted the dimension $\dim(V_\xi)$ of the representation, thus

$$\int_{\hat{G}} F(\xi)\ d\xi := \sum_{\xi \in \hat{G}} \dim(V_\xi) F(\xi).$$

Let $HS(V_\xi)$ be the space of linear transformations on V_ξ. We endow this space with the *Hilbert-Schmidt inner product*

$$\langle A, B \rangle_{HS(V_\xi)} := \mathrm{tr}_{V_\xi}(AB^*)$$

where we are using the normalised trace. Thus the Fourier transform is a map from $L^2(G)$ to the direct integral $\int_{\hat{G}} HS(V_\xi)\ d\xi$, which is the Hilbert space with inner product

$$\langle A, B \rangle_{\int_{\hat{G}} HS(V_\xi)\ d\xi} := \int_{\hat{G}} \langle A(\xi), B(\xi) \rangle_{HS(V_\xi)}\ d\xi.$$

It turns out the correct inverse Fourier transform is the following. Let $F \in \int_{\hat{G}} HS(V_\xi)\ d\xi$, then we define the function $\check{F} \in L^2(G)$ by

$$\check{F}(x) := \int_{\xi \in \hat{G}} \langle F(x), \rho_\xi(x) \rangle\ d\xi.$$

This is indeed the inverse Fourier transform:

Proposition 2.1. *Let $F \in \int_{\hat{G}} HS(V_\xi)\ d\xi$. Then $\widehat{\check{F}} = F$.*

Proof By linearity we may assume that F has only one non-zero component, e.g. $F(\xi) = A1_{\xi=\xi_0}$ for some $\xi_0 \in \hat{G}$ and $A \in HS(V_{\xi_0})$. In fact by linearity again we may assume that A is a rank one operator $Av = \langle v, a\rangle b$ for some unit vectors $a, b \in V_{\xi_0}$. Then

$$\check{F}(x) = \dim(V_{\xi_0})\mathrm{tr}_{V_{\xi_0}}(A\rho_{\xi_0}(x)^*) = \dim(V_{\xi_0})\langle \rho_{\xi_0}(x)^* b, a\rangle$$

and

$$\widehat{\check{F}}(\xi) = \dim(V_{\xi_0}) \int_G \langle \rho_{\xi_0}(x)^* b, a\rangle \rho_\xi(x)\ dx$$

and so for any $c, d \in V_\xi$

$$\langle \widehat{\check{F}}(\xi)c, d\rangle = \dim(V_{\xi_0}) \int_G \langle \rho_{\xi_0}(x)^* b, a\rangle \langle \rho_\xi(x)c, d\rangle\ dx$$

$$= \dim(V_{\xi_0})\langle \int_G \rho_\xi(x) T_{ca} \rho_{\xi_0}(x)^* b, d\rangle$$

where $T_{ca}(v) := \langle v, a\rangle c$. When $\xi \neq \xi_0$ the right-hand side vanishes thanks to Corollary 1.10. Thus it will suffice to show that

$$\langle c, a\rangle\langle b, d\rangle = \dim(V_{\xi_0})\langle \int_G \rho_{\xi_0}(x) T_{ca} \rho_{\xi_0}(x)^* b, d\rangle;$$

but this follows from Corollary 1.10 and the observation that $\mathrm{tr}_{V_{\xi_0}}(T_{ca}) = \langle c, a\rangle$. ■

Let $F \in \int_{\hat{G}} HS(V_\xi)\ d\xi$. From Fubini's theorem we have

$$\|\check{(F)}\|^2_{L^2(G)} = \langle \check{F}, \check{F}\rangle_{L^2(G)} = \int_{\hat{G}} \langle \widehat{\check{F}}, F\rangle_{HS(V_\xi)}\ d\xi$$

and so by the above proposition we have a preliminary Plancherel identity

$$\|\check{(F)}\|^2_{L^2(G)} = \int_{\hat{G}} \|F\|^2_{HS(V_\xi)}\ d\xi.$$

As a consequence, we see that the inverse Fourier transform is an isometry from $\int_{\hat{G}} HS(V_\xi)\ d\xi$ to $L^2(G)$, thus the forward Fourier transform is a co-isometry - a projection followed by a unitary transformation. But

Proposition 2.2. *The Fourier transform is injective.*

Proof Suppose we had $f \in L^2(G)$ such that $\hat{f} = 0$, thus $\int_G f(x)\rho(x)\ dx = 0$ whenever ρ is an irreducible representation. Taking direct sums, we see the same is true for arbitrary representations, and in particular for the regular representation. In other words, $f * g = 0$ for all $g \in L^2(G)$. Taking g to be a Dirac function we obtain $f = 0$, obtaining injectivity. ■

Thus the Fourier transform must in fact be a unitary transformation from $L^2(G)$ to $\int_{\hat{G}} HS(V_\xi)\ d\xi$. In particular we have the *Fourier inversion formula*

$$f(x) = \int_{\hat{G}} \langle \hat{f}(\xi), \rho_\xi(x)\rangle_{V_\xi}\ d\xi$$

and the *Plancherel identity*

$$\|f\|^2_{L^2(G)} = \int_{\hat{G}} \|\hat{f}(\xi)\|^2_{HS(V_\xi)}\ d\xi$$

from which one quickly deduces the *Parseval identity*

$$\langle f, g \rangle_{L^2(G)} = \int_{\hat{G}} \langle \hat{f}(\xi), \hat{g}(\xi) \rangle_{HS(V_\xi)} \, d\xi.$$

The regular representation is an action of G on $L^2(G)$, and is thus (via the Fourier transform) isomorphic to the obvious action of G on $\int_{\hat{G}} HS(V_\xi) \, d\xi$. It is not hard to see that the action of G on $HS(V_\xi)$ splits into the direct sum of $\dim(V_\xi)$ representations isomorphic to V_ξ. Thus the multiplicity of V_ξ in the regular representation is $\dim(V_\xi)$. Taking dimensions, we also observe the formula

$$\#G = \sum_{\xi \in \hat{G}} \dim(V_\xi)^2.$$

We remark that when G is abelian, this Fourier transform collapses to the previous one except for an annoying complex conjugation sign, which comes from a confluence of several mildly inconsistent notational conventions.

2.3. **Class functions and characters.** It is perhaps a little disconcerting to see the Fourier transform turn into an operator-valued function rather than a scalar-valued one. But as we saw, the noncommutativity of convolution leaves us little choice in this matter. However, things become simpler if one works in a reduced set of functions in which the effects of non-abelian-ness have been suppressed.

Definition 2.4. A *class function* is a function $f \in L^2(G)$ which is conjugation invariant, thus $f(y^{-1}xy) = f(x)$ for all $x, y \in G$. Equivalently, a class function is a function which is constant on each of the conjugacy classes of G. The space of all class functions is denoted $L^2(G)_G$.

Problem 2.5. Without using the Fourier transform, show that $L^2(G)_G$ is closed under convolution, and furthermore that convolution is commutative in this space.

This leads one to hope that a scalar theory is now possible. And indeed it is:

Lemma 2.6. *A function $f \in L^2(G)$ is a class function if and only if $\hat{f}(\xi)$ is a multiple of the identity id_{V_ξ} for each $\xi \in \hat{G}$.*

Proof Suppose first that f is a class function, then $f(y^{-1}xy) = f(x)$ for all y. Taking Fourier transforms in x, we conclude that $\rho_{V_\xi}(y)\hat{f}(\xi)\rho_{V_\xi}(y)^{-1} = \hat{f}(\xi)$ for all $\xi \in \hat{G}$ and $y \in G$. Applying Corollary 1.10 we conclude that $\hat{f}(\xi)$ is a multiple of the identity. Conversely, if $\hat{f}(\xi) = c_\xi \mathrm{id}_{V_\xi}$ for some complex scalars c_ξ then by the inversion formula

$$f(x) = \int_G c_\xi \overline{\mathrm{tr}_{V_\xi}(\rho_\xi(x))} \, d\xi.$$

Each function $\mathrm{tr}_{V_\xi}(\rho_\xi(x))$ can be seen to be a class function, and so f is a class function as desired. ∎

Note that this lemma immediately implies Problem 2.5. As a corollary, we also see that the number of conjugacy classes of G (which is the dimension of $L^2(G)_G$) is also equal to $\#\hat{G}$.

For each representation $\rho_V : G \to U(V)$, define the associated *character* $\chi_V(x) := \mathrm{tr}_V(\rho_V(x))$; thus for instance $\chi_V(0)$ is the dimension of V. As observed in the proof of the above lemma, characters are class functions. In fact

Proposition 2.7. *The characters $\{\chi_\xi : \xi \in \hat{G}\}$ form an orthonormal basis of $L^2(G)_G$.*

We remark that the basis of characters, written out as a function of the conjugacy classes, is known as the *character table* of G, and plays a fundamental role in the structural theory of finite groups.

Proof From the proof of the previous lemma we see that the characters (or more precisely, the conjugates of these characters) span $L^2(G)_G$. Also, observe that $\overline{\chi_{\xi_0}}$ is the inverse Fourier transform of $\frac{1}{\dim(V_\xi)}\mathrm{id}_{V_\xi}$. The orthonormality then follows from Parseval. ∎

Observe that $\chi_{V\oplus W} = \chi_V + \chi_W$ and $\chi_{V\otimes W} = \chi_V \chi_W$. Also, isomorphic representations have the same character. As a consequence of this, we see that for any representation V we have $\chi_V = \sum_{\xi\in\hat{G}} m_\xi \chi_{V_\xi}$, where m_ξ is the multiplicity of ξ in V; by orthonormality we thus conclude that $m_\xi = \langle \chi_V, \chi_{V_\xi}\rangle$. In particular the isomorphism class of the representation V is determined entirely by the character V. From the orthonormality we have

$$\int_G |\chi_V|^2 \, dx = \sum_{\xi\in\hat{G}} m_\xi^2. \tag{1}$$

In particular we see that V is irreducible if and only if $\int_G |\chi_V|^2 \, dx = 1$; this gives a convenient way to test for irreducibility.

Let us now illustrate some of this theory with the simplest example of a non-abelian finite group, namely the permutation group S_3 of three elements $\{1,2,3\}$. This group has three conjugacy classes: the identity $C_1 = \{\mathrm{id}\}$, the class $C_2 = \{(12),(23),(31)\}$ of transpositions, and the class $C_3 = \{(123),(231)\}$ of cycles. Thus we expect three irreducible representations up to isomorphism. One of them is the trivial representation on $\mathbf{C}$, whose character equals 1 on C_1, C_2, C_3. Another is the alternating representation on $\mathbf{C}$, using the sign character sgn; this character equals -1 on C_2 and $+1$ on C_1, C_3. Finally, we have the standard representation of S_3 on the plane $\{(x_1,x_2,x_3) \in \mathbf{C}^3 : x_1+x_2+x_3 = 0\}$, on which S_3 acts by coordinate permutation; the character here equals 2 on C_1, 0 on C_2, and -1 on C_3. One can verify that the character table is indeed orthonormal as claimed.

3. Expansion in $SL_n(F_q)$

We now use the above theory to demonstrate a certain "quasirandomness" property of the special linear group $G := SL_n(F_q)$, defined as the group of $n \times n$ matrices with determinant 1 and coefficients in a finite field F_q of q elements (where q is a prime, or a power of a prime). To avoid some technicalities we assume that q as being odd. We think of $n \geq 2$ as being fixed (e.g. $n = 2$) and q being large. It is not hard to see that the cardinality of G is $\sim_n q^{n^2-1}$; more precise formulae are available but we do not need them here.

Characterising all the irreducible representations of this group is a non-trivial task, involving a fair amount of combinatorial and number-theoretic ingenuity. We will however content outselves with the following simple result of Frobenius.

Lemma 3.1. *Let V be a non-trivial irreducible representation of $G = SL_n(F_q)$. Then V has dimension at least $\frac{q-1}{2}$.*

Contrast this with the irreducible representations of an abelian group, which are always one dimensional. Thus this lemma indicates that the special linear groups are highly non-abelian in some sense, especially when q is large. The bound $\frac{q-1}{2}$ is sharp, but we will not establish that here.

Proof We shall analyse V using one-dimensional abelian subgroup of G, e.g. the group

$$H := \{g_x : x \in F_q\}$$

where g_x is the matrix

$$g_x := \begin{pmatrix} 1 & x & 0 & \dots \\ 0 & 1 & 0 & \dots \\ 0 & 0 & 1 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}.$$

This group is isomorphic to the *additive* group F_q. As V is a representation of G, it is also a representation of H, and we can restrict the character $\chi_V : G \to \mathbf{C}$ to H and view it as a character $\chi_V|_H$ on H. We can split V into irreducible representations of H (which are all one-dimensional), and thus split $\chi_V|_H$ as the sum of multiplicative characters on H; the number of summands is then the dimension of V. Thus it will suffice to show that the Fourier transform of $\chi_V|_H$ has support of size at least $\frac{q-1}{2}$.

Now observe that for any non-zero $a \in F_q$, g_x and g_{a^2x} are conjugate to each other in $SL_n(F_q)$. Thus $\chi_V|_H$, when viewed as a function on F_q, is invariant under dilations by squares a^2; thus the Fourier transform has the same property. Thus if the Fourier transform is nonzero in even just a single non-zero frequency ξ, it is automatically non-zero at the $\frac{q-1}{2}$ frequencies $\{a^{-2}\xi : a \in F_q \backslash \{0\}\}$. Thus we are done unless the Fourier transform vanishes at all non-zero frequencies, i.e. $\chi_V|_H$ is constant. But this implies that $\rho_V(g_x) = \mathrm{id}_V$ for all $g_x \in H$. Conjugating this, we then conclude that $\rho_V(g)$ is the identity for any g in any of the conjugates of H. But these are easily seen to generate all of $SL_n(F_q)$ (indeed, it is not hard to see that they generate both the upper-triangular and lower-triangular unipotent

matrices, so the group generated is closed under both row and column operators, which makes it quite easy to generate arbitrary elements of $SL_n(F_q)$). Thus V is trivial, a contradiction. ∎

Corollary 3.2 (Kunze-Stein type estimate)**.** *Let $G = SL_2(F_q)$, and let $f, g \in L^2(G)$. If at least one of f, g has mean zero, then*

$$\|f * g\|_{L^2(G)} \lesssim q^{-1/2} \|f\|_{L^2(G)} \|g\|_{L^2(G)}.$$

Of course, the point here is the gain of $q^{-1/2}$ over what one might expect from Young's inequality. Without the mean zero hypothesis, the estimate fails, as can be seen by taking $f = g = 1$.

Proof Applying Plancherel's theorem, the left-hand side is

$$(\sum_{\xi \in \hat{G}} \dim(V_\xi) \|\hat{f}(\xi)\hat{g}(\xi)\|^2_{HS(V_\xi)})^{1/2}$$

and the right-hand side is

$$q^{-1/2} (\sum_{\xi \in \hat{G}} \dim(V_\xi) \|\hat{f}(\xi)\|^2_{HS(V_\xi)})^{1/2} (\sum_{\xi \in \hat{G}} \dim(V_\xi) \|\hat{g}(\xi)\|^2_{HS(V_\xi)})^{1/2}.$$

Since f or g has mean zero, we see that the contribution of the trivial representation $\xi = 0$ to the LHS vanishes. Thus we can restrict attention to non-trivial representations. It thus suffices to show that

$$\dim(V_\xi) \|\hat{f}(\xi)\hat{g}(\xi)\|^2_{HS(V_\xi)} \lesssim q^{-1} \dim(V_\xi) \|\hat{f}(\xi)\|^2_{HS(V_\xi)} \dim(V_\xi) \|\hat{g}(\xi)\|^2_{HS(V_\xi)}$$

for all $\xi \neq 0$l in view of the previous lemma, it thus suffices to establish the algebra estimate

$$\|AB\|_{HS(V)} \leq \|A\|_{HS(V)} \|B\|_{HS(V)}.$$

This in turn follows from two easy estimates. The first is that the Hilbert-Schmidt norm controls the operator norm:

$$\|A\|_{V \to V} \leq \|A\|_{HS(V)}.$$

This is easily seen by working in a orthonormal basis and using duality and Cauchy-Schwarz. The second is that the Hilbert-Schmidt space is an ideal in the space of bounded operators:

$$\|AB\|_{HS(V)} \leq \|A\|_{V \to V} \|B\|_{HS(V)}.$$

This is immediate from the identity

$$\|B\|_{HS(V)} = (\sum_i \|Be_i\|^2_V)^{1/2}$$

where e_i ranges over an arbitrary orthonormal basis of V. ∎

This leads to the following corollary. Let A be any symmetric subset of G (thus $A = A^{-1}$) not containing the identity, and consider the *Cayley graph* (V, E) whose vertex set V is just the group G, and the edge set consists of all pairs x, y such that $xy^{-1} \in A$. Thus each vertex is connected by $\#A$ edges to other vertices. Let B, C be any other subsets of G, and consider the quantity

$$E(B, C) := \{(b, c) \in B \times C : (b, c) \in E\},$$

i.e. the number of ways B and C are connected together by the Cayley graph. Since the "edge density" of G is roughly $\#A/\#G$, one would expect that $E(B,C)$ would be roughly $\#A\#B\#C/\#G$ if $E(B,C)$ is sufficiently "quasirandom". This is indeed the case (an observation of Gowers):

Lemma 3.3. *We have*

$$E(B,C) = \frac{\#A\#B\#C}{\#G} + O(q^{-1/2}(\#A\#B\#C\#G)^{1/2}).$$

Note that the error term is dominated by the main term if A, B, C are fairly dense subsets of G. This statement is an assertion that all dense Cayley graphs in $SL_n(F_q)$ are rather quasirandom in nature.

Proof We can write the left-hand side as

$$(\#G)^2 \int_G 1_A * 1_B(x) 1_C(x)\ dx.$$

We write $1_B = (1_B - \int_G 1_B) + (\int_G 1_B)$. The contribution of the second term can easily be computed to be exactly $\frac{\#A\#B\#C}{\#G}$. The contribution of the first term can be bounded by $O(q^{-1/2}(\#A\#B\#C\#G)^{1/2})$ using the previous lemma and Cauchy-Schwarz. ■

Further estimates in this direction have been obtained. One recent (and rather non-trivial) result of Bourgain and Gamburd is the following: if p is prime, and $S = O(1)$ is a symmetric set in G of girth $\gtrsim \log p$ (i.e. the smallest non-trivial word in S which multiplies to the identity has length $\gtrsim \log p$) then the Cayley graph associated to S is an expander, i.e. we have $|A \cdot S| \geq (1+\varepsilon)|A|$ for all $|A| \leq |G|/2$ and some $\varepsilon > 0$ depending only on S. This has applications to random number generation, property T, the Ramanujan conjectures, and a wide variety of other interesting mathematical topics!

4. A brief discussion of compact Lie groups

It is highly non-trivial to extend the representation theory - especially the infinite dimensional representation theory - of finite groups to arbitrary infinite groups, even after restricting to unitary representations of Lie groups. For instance, consider the Heisenberg group

$$H = \{(x,\xi,t) : x,\xi \in \mathbf{R}; t \in \mathbf{R}/\mathbf{Z}\}$$

with multiplication law

$$(x,\xi,t)\cdot(x',\xi',t') := (x+x', \xi+\xi', t+t'-x\xi').$$

This acts on unitarily on the infinite dimensional space $L^2(\mathbf{R})$ by the translation-modulation action

$$\rho(x,\xi,t)f(y) := e^{2\pi i t} e^{2\pi i y\xi} f(y-x)$$

but this action has no finite-dimensional subrepresentations. (The problem here is that the commutator of the infinitesimal translation $(dx,0,0)$ and the infinitesimal

modulation $(0, dt, 0)$ is a multiple of the identity, but in finite dimensions commutators must have zero trace.)

However, the situation is much more manageable when restricting attention to *compact* Lie groups, such as the rotation group $O(n)$ or the unitary group $U(n)$. Here, it turns out that much of the finite theory still applies, for instance every representation can be viewed as a unitary representation, that indecomposable representations are irreducible, that irreducible representations can be viewed as subrepresentations of $L^2(G)$, one has a good character theory, and so forth. One particularly nice feature of compact Lie groups is that they are *unimodular* - the left-Haar measure dx is also the right-Haar measure. (This is because right-translates of dx are still left-invariant, and assign the same total mass to the whole group, and so must equal dx.)

The main technical difficulty is to ensure that the irreducible representations of $L^2(G)$ are all finite dimensional. One way to do this is to use the *Casimir operator* Ω on G, which is an analogue of the Laplacian, but where the metric is replaced instead by the *Killing form* $B(X, Y) = \mathrm{tr}(\mathrm{ad}(X)\mathrm{ad}(Y))$. The main point is that the Casimir operator commutes with all the vector fields in the Lie algebra, and as a consequence the eigenspaces of the Casimir operator are subrepresentations of the regular representation. On the other hand, the Casimir is a positive-definite operator with compact resolvents, and standard spectral theory methods show that the spectrum is purely discrete and that all eigenspaces are finite-dimensional. From these facts it is possible to split the regular representation into a countable direct sum of finite-dimensional irreducible representations. We will however not dwell on these matters here.

5. Exercises

- Q1. Let $\rho : G \to GL(V)$ be a non-unitary finite-dimensional representation of a finite group G on a vector space V. Show that there exists a Hilbert space structure on V on which the action becomes unitary. (Hint: start with an arbitrary Hilbert space structure and average it.)
- Q2. Let $\rho_V : G \to U(V)$ and $\rho_W : G \to U(W)$ be two irreducible representations of a finite group G, and let $\mathrm{Hom}(V, W)$ be the space of linear transformations from V to W. The product group $G \times G$ acts on this space by the formula
$$\rho_{\mathrm{Hom}(V,W)}(g, h)(T) := \rho_W(h) T \rho_V(g)^{-1}.$$
Show that this representation of $G \times G$ is irreducible.
- Q3. Let H be a subgroup of a finite group G, and let V be an irreducible representation of G. Show that if V is viewed instead as a representation of H, then V decomposes into at most $\#G/\#H$ irreducible representations. (Hint: use (1).)

Department of Mathematics, UCLA, Los Angeles CA 90095-1555

E-mail address: tao@math.ucla.edu