Small Representations of Finite Classical Groups

Roger Howe

Yale University

Singapore , March 2016

Problems from Group Theory

Note: this note summarizes joint work with Shamgar Gurevich.

Several interesting problems in group theory can be approached by using characters of irreducible representations.

- 1. Convergence of random walks:

  Given a probability measure on a simple group G (i.e.,, an element $\mu = \sum_{g \in G} c_g g$ in the group algebra, with coefficients c_g satisfying $c_g \geq 0$ and $\sum_{g \in G} c_g = 1$) at what rate do the convolution powers of μ approach the constant measure?

  In particular, if μ is a normalized conjugacy class C: $c_g = \frac{1}{\#(C)}$ if $g \in C$, and $c_g = 0$ otherwise?

Problems from Group Theory

Note: this note summarizes joint work with Shamgar Gurevich.

Several interesting problems in group theory can be approached by using characters of irreducible representations.

- 1. Convergence of random walks:

  Given a probability measure on a simple group G (i.e.,, an element $\mu = \sum_{g \in G} c_g g$ in the group algebra, with coefficients c_g satisfying $c_g \geq 0$ and $\sum_{g \in G} c_g = 1$) at what rate do the convolution powers of μ approach the constant measure?

  In particular, if μ is a normalized conjugacy class C: $c_g = \frac{1}{\#(C)}$ if $g \in C$, and $c_g = 0$ otherwise?

- 2. Ore Conjecture: Is every element in a finite simple group equal to a commutator?

Problems from Group Theory

Note: this note summarizes joint work with Shamgar Gurevich.

Several interesting problems in group theory can be approached by using characters of irreducible representations.

- 1. Convergence of random walks:

  Given a probability measure on a simple group G (i.e.,, an element $\mu = \sum_{g \in G} c_g g$ in the group algebra, with coefficients c_g satisfying $c_g \geq 0$ and $\sum_{g \in G} c_g = 1$) at what rate do the convolution powers of μ approach the constant measure?

  In particular, if μ is a normalized conjugacy class C: $c_g = \frac{1}{\#(C)}$ if $g \in C$, and $c_g = 0$ otherwise?

- 2. Ore Conjecture: Is every element in a finite simple group equal to a commutator?

- 3. Thompson Conjecture: In a finite simple group G, is there a conjugacy class C such that $C^2 = G$?

Given a finite group G, let

$\hat{G}$ = the set of equivalence classes of irreducible representations of G (over $\mathbf{C}$).

1 = trivial representation of G,

$\hat{G} - \{1\}$ = *non-trivial* representations of G.
For a given representation ρ, the *character* χ_ρ is the function

$$\chi_\rho(g) = \mathit{trace}(\rho(g)).$$

In particular, if e_G is the identity element of G, then

$$\chi_\rho(e_G) = \dim \rho$$

is the dimension of ρ.

With this terminology, the questions stated above can be translated into statements about characters, as follows.

Random Walks

In question 1, the rate of convergence is controlled by

$$\max_{1 \neq \rho \in \hat{G}} \frac{\chi_\rho(C)}{\dim \rho},$$

over all non-trivial ρ in $\hat{G}$.

Counting Commutators

In question 2, the number of ways to express g as a commutator of two other elements of G is

$$\#(G)\left(\sum_{\rho\in\hat{G}}\frac{\chi_\rho(g)}{\dim\rho}\right).$$

Squaring a Conjugacy Class

In question 3, the number of ways to express a given g in G as the product of two elements in the conjugacy class C is

$$\frac{\#(C)^2}{\#(G)} \left(\sum_{\rho \in \hat{G}} \chi_\rho(C)^2 \frac{\chi_\rho(g)}{\dim \rho} \right).$$

In the second and third expressions, there is always a positive term coming from the trivial representation. Thus, one would like to have some control on the sum over non-trivial characters to show that the full sum is non-zero.

In particular, if one could show that the terms involving non-trivial ρ were small enough (and in addition, perhaps, show that some terms cancel each other, or nearly so), then one could conclude that the sum was non-zero.

Character Ratios

This leads to focus on the terms

$$\frac{\chi_\rho(C)}{\dim \rho} \quad - \quad \textit{the character ratios}.$$

One would like to know that the character ratios are small.
The part that is making them small is the denominator $\dim \rho$.
So one would like to know about the behavior of the dimensions of the irreducible representations.
In particular:

Describe the representations of G with "small" dimensions.

Of course, "small" here is relative to some "typical" or "large" dimension of representations of G.

II. Example: symplectic groups

We will discuss the question from I for a particular case, the symplectic groups over finite fields. Recall that

i) the symplectic groups are one class of (almost) simple algebraic groups; and

- ii) the finite groups of Lie type, equivalently, the simple algebraic groups over finite fields, constitute the large majority of finite simple groups; and

II. Example: symplectic groups

We will discuss the question from I for a particular case, the symplectic groups over finite fields. Recall that

i) the symplectic groups are one class of (almost) simple algebraic groups; and

- ii) the finite groups of Lie type, equivalently, the simple algebraic groups over finite fields, constitute the large majority of finite simple groups; and
- iii) the classical groups constitute the large majority of groups of Lie type; and

II. Example: symplectic groups

We will discuss the question from I for a particular case, the symplectic groups over finite fields. Recall that

i) the symplectic groups are one class of (almost) simple algebraic groups; and

- ii) the finite groups of Lie type, equivalently, the simple algebraic groups over finite fields, constitute the large majority of finite simple groups; and
- iii) the classical groups constitute the large majority of groups of Lie type; and
- iv) the symplectic groups are one of four (or three, depending how you count) series of classical groups.

II. Example: symplectic groups

We will discuss the question from I for a particular case, the symplectic groups over finite fields. Recall that

i) the symplectic groups are one class of (almost) simple algebraic groups; and

- ii) the finite groups of Lie type, equivalently, the simple algebraic groups over finite fields, constitute the large majority of finite simple groups; and
- iii) the classical groups constitute the large majority of groups of Lie type; and
- iv) the symplectic groups are one of four (or three, depending how you count) series of classical groups.
- These facts together suggest that knowing the behavior of symplectic groups with respect to this question is both significant for the overall answer, and suggestive of behavior for other classes.

Symplectic groups

So, let $\mathbf{F}_q$ be the finite field with q elements, and let $W = (\mathbf{F}_q)^{2n}$ be equipped with the skew-symmetric, non-degenerate, bilinear form (aka *symplectic* form)

$$< \vec{w}, \vec{w}' >= \sum_{j=1}^{n} w_j w'_{n+j} - w_{n+j} w'_j,$$

where the w_j are the entries of the element $\vec{w} \in W$, and similarly for the w'_j. Then

$$Sp(W) = Sp_{2n}(\mathbf{F}_q)$$

is the subgroup of $GL(W)$ consisting of elements that preserve the pairing $< \, , \, >$.

Lagrangian subspaces

The subspace X of W, consisting of $\vec{x}$ for which $x_j = 0$ for $j > n$, is *totally isotropic* for $< \, , \, >$; that is the restriction of $< \, , \, >$ is identically zero; and X is maximal with respect to being totally isotropic. It is sometimes called a *Lagrangian* subspace.

The subgroup P_X of $Sp(W)$ that preserves X is called the *Siegel parabolic*.

The Siegel Parabolic subgroup

P_X has the form

$$P_X \simeq GL(X) \cdot U_X$$

Here the group U_X is normal in P_X and is abelian.
More precisely,

$$U_X \simeq S_n^2(\mathbf{F}_q),$$

$$= \text{the space of } n \times n \text{ symmetric matrices over } \mathbf{F}_q.$$

Likewise,

$$\hat{U}_X = \text{Pontrjagin dual of } U_X,$$

is isomorphic to S_n^2, under the mapping

$$\psi_A(B) = \chi(traceAB)$$

for symmetric matrices A, B, and a fixed character χ of $\mathbf{F}_q$.

The action of $GL(X)$ on these groups is the standard action of GL_n on symmetric matrices.

Restricting representations to the Siegel unipotent

Consider a representation ρ of $Sp(W)$, and look at the restriction to U_X. It will decompose as a sum of characters, with certain multiplicities:

$$\rho_{|U_X} \simeq \sum_{B \in S_n^2} m_B \psi_B.$$

Since

$$\rho_{|U_X} = \left(\rho_{|P_X}\right)_{|U_X},$$

that is,
the restriction of ρ in $\widehat{Sp}$ to U_X can be thought of as the restriction to U_X of the restriction of ρ to P_X,

the U_X spectrum of $\rho_{|U_X}$ must be invariant under the action of $GL(X)$. That is, $m_B = m_{B'}$ if B and B' define equivalent symmetric bilinear forms on X.

Isomorphism classes of symmetric bilinear forms

The first major invariant of a symmetric bilinear form is its rank.

It is well known that, over finite fields, there are just 2 isomorphism classes of bilinear forms of a given rank k.

We denote these classes by $\mathcal{O}_{k+}$ and $\mathcal{O}_{k-}$;
or we will denote the pair of them, or whichever one is relevant in a given context as $\mathcal{O}_{k\pm}$.

Rank of characters of U_X

If B is a form of rank k, we will also say that the associated character ψ_B has rank k. We may also refer to the character as being of type $+$ or type $-$, according to the type of B.

With this notation, we can organize the description of $\rho_{|U_X}$ according to the ranks, and also according to the isomorphism classes of the associated forms:

$$\rho_{|U_X} \simeq \sum_k \sum_{\pm} m_{k\pm} \left(\sum_{B \in \mathcal{O}_{k\pm}} \psi_B \right).$$

This formula implies that the dimension of ρ must be a sum of the cardinalities $\#(\mathcal{O}_{k\pm})$ of the isomorphism classes of symmetric bilinear forms.

Sizes of isomorphism classes of bilinear forms

It is easy to give a formula for these cardinalities. We have

$$\#(\mathcal{O}_{k\pm}) = \Gamma_{nk}\frac{\#(GL_k)}{\#(O_{k\pm})}.$$

In this formula,

Γ_{nk} indicates the cardinality of the Grassmann variety of k dimensional subspaces of $(\mathbf{F}_q)^n$; and

$O_{k\pm}$ indicates the isometry group of a non-degenerate form of type $\pm$ on $(\mathbf{F}_q)^k$.

From the formula,

$$\#(\mathcal{O}_{k\pm}) \simeq \tfrac{1}{2} q^{k(n-\frac{k-1}{2})}.$$

In particular, the smallest orbits are those of rank one forms. These have size

$$\frac{q^n - 1}{2}.$$

The smallest possible representation

It follows from this discussion that
the smallest possible dimension of a non-trivial representation of $Sp(W)$ would be

$$\frac{q^n - 1}{2}.$$

- Such a representation would contain each rank one character of one type, and nothing else.

  Since U_X is such a small subgroup of $Sp(W)$, it is unclear whether to expect such a representation to exist. In particular, it would be irreducible already on P_X, and it would be the smallest possible faithful representation of P_X.

- Such a representation would contain each rank one character of one type, and nothing else.

  Since U_X is such a small subgroup of $Sp(W)$, it is unclear whether to expect such a representation to exist. In particular, it would be irreducible already on P_X, and it would be the smallest possible faithful representation of P_X.
- It turns out, however, that it does exist; in fact, there are two.

  Proposition: There are two irreducible representations of $Sp(W)$ of dimension $\frac{q^n-1}{2}$, one containing either one of the two rank one $GL(X)$ orbits in $\hat{U}_X$

The next smallest possible representation

What is the next largest possible dimension? Well, one more - the U_X spectrum could include a rank one orbit, and a trivial representation. It turns out that these also exist.

Proposition: There are two irreducible representations of $Sp(W)$ of dimension $\frac{q^n-1}{2} + 1 = \frac{q^n+1}{2}$, one whose U_X-spectrum contains one of the rank one orbits in $\hat{U}_X$.

Rank of representations of Sp

These results, plus considerations of tensor products, tell us that, for any orbit $\mathcal{O}_{k\pm}$ in $\hat{U}_X$, there will be representations of $Sp(W)$ whose U_X spectrum contains the given orbit, together with orbits of smaller rank. Since the size of the orbits $\mathcal{O}_{k\pm}$ is increasing with k, representations whose U_X spectrum is concentrated on orbits of smaller rank can be expected to have smaller dimensions. This motivates the following definition.

Definition: a) A representation ρ in $\widehat{Sp(W)}$ is of *rank k* iff the restriction $\rho_{|U_X}$ contains characters of rank k, but of no higher rank.
b) If a representation ρ in $\widehat{Sp(W)}$ of rank k contains characters of type $\mathcal{O}_{k+}$, but not of type $\mathcal{O}_{k-}$, then we say that ρ is of type $\mathcal{O}_{k+}$; and likewise with $+$ and $-$ switched.

Heisenberg groups

III. The Heisenberg group and the Oscillator Representation

Where do the smallest representations of $Sp(W)$ come from?

They can be found by considering the *Heisenberg group* associated to W. This is the group

$$H(W) = W \oplus \mathbf{F}_q,$$

with group law defined by

$$(\vec{w}, z) \cdot (\vec{w}', z') = (\vec{w} + \vec{w}', z + z' + \frac{1}{2} < \vec{w}, \vec{w}' >).$$

This is a two-step nilpotent group.

$$\text{Center of } H(W) = \text{commutator subgroup of } H(W) =$$

$$Z = \{(0, z)\}.$$

The commutator operation in $H(W)$ induces a skew-symmetric bilinear form on $H(W)/Z \simeq W$, that coincides with the original symplectic form.

The group $H(W)$ is the analog over a finite field of the Lie group associated to the Canonical Commutation Relations (CCR) of W. Heisenberg, of Uncertainty Relation fame.

When $q = p$ is prime, Heisenberg groups are important in group theory and were known as *extra special p groups*.

The representation theory of Heisenberg groups is very simple. For any representation ρ in $\widehat{H(W)}$, the center Z will act by scalars:

$$\rho((0,z)) = \chi_\rho(z)I,$$

where I is the identity operator, and $\chi_\rho = \chi \in \hat{Z}$ is a character of Z, called the *central character* of ρ.

If $\chi_\rho = 1$, then ρ factors to $H(W)/Z \simeq W$, which is abelian, so ρ itself is a character of W.

The case of non-trivial χ_ρ is described by

Representations of $H(W)$

Stone-von Neumann-Mackey Theorem: Up to equivalence, there is a unique irreducible representation ρ_χ with given non-trivial central character χ in $\hat{Z} - \{1\}$.

Remarks: a) There are many ways to realize the presentation ρ explicitly.
b) It particular, it can be constructed as an induced representation from any character extending χ to any maximal abelian subgroup of $H(W)$.

c) The inverse image in $H(W)$ of any maximal isotropic subspace (e.g., X) of W will be a maximal abelian subgroup of $H(W)$.

Constructing the oscillator representation I

Definition of $H(W)$, $\Rightarrow$

a) the action of $Sp(W)$ on W lifts to an action of $H(W)$ by automorphisms,

b) leaving the center point wise fixed.

Hence
c) the induced action of $Sp(W)$ on $\widehat{H(W)}$ will leave each ρ_χ fixed, for any character $\chi \in \hat{Z} - \{1\}$.

Constructing the oscillator representation II

Hence
d) for each g in $Sp(W)$, there is an operator $\omega(g)$ on ρ_χ (more properly, on the vector space where $\rho\chi$ is realized), such that

$$\omega(g)\rho_\chi(h)\omega(g)^{-1} = \rho_\chi(g(h).$$

The operator $\omega(g)$ is defined up to scalar multiples. This implies that

$$\omega(g)\omega(g') = \alpha(g, g')\omega(gg'),$$

where $\alpha(g, g')$ is an appropriate complex number, of absolute value 1.

Constructing the oscillator representation III

For finite fields, it is known that this mapping can be lifted to a genuine representation.
That is, the operators $\omega(g)$ can be chosen so that $\alpha(g, g') \equiv 1$ for all g, g' in $Sp(W)$.

We call it the *oscillator representation*

(aka: *Weil representation*).

In fact, we should have used the notation ω_χ, because the oscillator representation depends on the character χ.

However, this dependence is weak.
There are only two possible oscillator representations.

We have $\omega_\chi \simeq \omega_{\chi'}$ if $\chi'((0, z)) = \chi(0, s^2 z)$ for some s in $(\mathbf{F}_q)^\times$.

The structure of the representations ω_χ is known explicitly. In particular, each ω_χ is of rank 1, and more precisely of type $\mathbf{O}_{1\pm}$, for an appropriate choice of $\pm$. That is, one of the ω_χ is of type $\mathbf{O}_{1+}$, and the other is of type $\mathbf{O}_{1-}$

Theorem: The representation ω_χ decomposes into two irreducible pieces, of dimensions $\frac{q^n \pm 1}{2}$.

Thus, the pair of representations ω_χ give concrete realizations of the irreducible representations of $Sp(W)$ of smallest possible dimensions.

IV. Representations of Rank k

We can use the oscillator representation also to construct representations of any rank k up to n. This is essentially by taking tensor products. However, a slightly different point of view gives us more insight into the nature of the rank k representations when $k < n$.

Let

$$U \simeq (\mathbf{F}_q)^k$$

be a vector space of dimension k over $\mathbf{F}_q$,
Let β be an inner product (non-degenerate, symmetric bilinear form) on U.

On the tensor product $W \otimes U$, the tensor product

$$< , > \otimes \beta$$

of the forms $< , >$ and β defines a symplectic form.

The isometry groups $Sp(W)$ and O_β of the forms $< \, , \, >$ and β act on $W \otimes U$
via their actions on the first and second factors respectively.
Both actions will preserve the form $< \, , \, > \otimes \beta$.
The action of $Sp(W)$ commutes with O_β, and vice versa.

Thus, we get an embedding

$$Sp(W) \times O_\beta \quad \hookrightarrow \quad Sp(W \otimes U).$$

The two factors $Sp(W)$ and O_β commute with each other.

In fact, each is the full centralizer of the other inside $Sp(W \otimes U)$.

Thus, the pair $(Sp(W), O_\beta)$ form what has been called a *dual pair* of subgroups of $Sp(W \otimes U)$.

Consider the restriction of the oscillator representation $\omega_{W\otimes U}$ to the product $Sp(W) \times O_\beta$.
(Note: we suppress the dependence of ω on the central character χ, but we record which symplectic group it belongs to.)

We decompose this restriction into isotypic components for O_β:

$$(\omega_{W\otimes U})_{|Sp(W)\cdot O_\beta} \simeq \sum_{\tau\in\hat{O}_\beta} \Theta(\tau)\otimes\tau,$$

where $\Theta(\tau)$ is a (not necessarily irreducible) representation of $Sp(W)$.

Although the factors $\Theta(\tau)$ will in general not be irreducible, we can say something about how they decompose.

- **Theorem**: a) As a representation of $Sp(W)$, $\omega_{W\otimes U}$ is of rank k, and of type β.

Although the factors $\Theta(\tau)$ will in general not be irreducible, we can say something about how they decompose.

- **Theorem**: a) As a representation of $Sp(W)$, $\omega_{W\otimes U}$ is of rank k, and of type β.
- b) For each τ in $\hat{O}_\beta$, the representation $\Theta(\tau)$ contains a unique irreducible constituent $\eta(\tau)$ of rank k; all other constituents have rank less than k.

Although the factors $\Theta(\tau)$ will in general not be irreducible, we can say something about how they decompose.

- **Theorem**: a) As a representation of $Sp(W)$, $\omega_{W\otimes U}$ is of rank k, and of type β.
- b) For each τ in $\hat{O}_\beta$, the representation $\Theta(\tau)$ contains a unique irreducible constituent $\eta(\tau)$ of rank k; all other constituents have rank less than k.
- c) The mapping $\tau \to \eta(\tau)$ gives an embedding

  $$\eta : \hat{O}_\beta \to \widehat{Sp(W)}_{k\beta} \subset \widehat{Sp(W)}_k.$$

  where
  $\widehat{Sp(W)}_k = \{$irreducible representations of $Sp(W)$ of rank k $\}$,
  and
  $\widehat{Sp(W)}_{k\beta} \subset \widehat{Sp(W)}_k$ of representations of type β.

Although the factors $\Theta(\tau)$ will in general not be irreducible, we can say something about how they decompose.

- **Theorem**: a) As a representation of $Sp(W)$, $\omega_{W\otimes U}$ is of rank k, and of type β.
- b) For each τ in $\hat{O}_\beta$, the representation $\Theta(\tau)$ contains a unique irreducible constituent $\eta(\tau)$ of rank k; all other constituents have rank less than k.
- c) The mapping $\tau \to \eta(\tau)$ gives an embedding

  $$\eta : \hat{O}_\beta \to \widehat{Sp(W)}_{k\beta} \subset \widehat{Sp(W)}_k.$$

  where
  $\widehat{Sp(W)}_k = \{$irreducible representations of $Sp(W)$ of rank k $\}$,
  and
  $\widehat{Sp(W)}_{k\beta} \subset \widehat{Sp(W)}_k$ of representations of type β.
- d) The multiplicity of the orbit $\mathcal{O}_\beta$ in $\eta(\tau)_{|U_W}$ is $\dim \tau$.

It also seems that this construction should produce all of $\widehat{Sp(W)}_k$.

Conjecture: We have

$$\widehat{Sp(W)}_k = \eta(\hat{O}_{\beta+}) \ \cup \ \eta(\hat{O}_{\beta-}),$$

where $\beta+$ and $\beta-$ represent the two isomorphism classes of inner products of rank k.

We have a partial argument for this conjecture, and it is supported by numerical calculations.

Remark: Some instances of the Theorem were observed by Aubert-Przebinda and by J. Epequin.

Example: $Sp_6(F_5)$

Dimensions:

Rank 0: 1

Rank 1: 62, 62, 63, 63

Rank 2: 1240, 1302, 1302, 1365, 1890, 1953, 1953,
2015, 2604, 2604, 3906

Rank 3: 6501, . . .

Principal series: 3, 656,016

Example: $Sp_6(F_3)$

Remarks/Questions:

- a) What does the map η look like in Lusztig parameters?

Example: $Sp_6(F_3)$

Remarks/Questions:

- a) What does the map η look like in Lusztig parameters?
- bi) A similar result holds for local fields.

  bii) How do these maps look in Langlands-Vogan parameters?

  biii) Can "rank k" be characterized simply in terms of wave front sets, or similar invariants?

Example: $Sp_6(F_3)$

Remarks/Questions:

- a) What does the map η look like in Lusztig parameters?
- bi) A similar result holds for local fields.

  bii) How do these maps look in Langlands-Vogan parameters?

  biii) Can "rank k" be characterized simply in terms of wave front sets, or similar invariants?
- ci) Assuming the conjecture, one can define "rank k" also for $n \leq k \leq 2n$.

  cii) How to describe rank k representations intrinsically for this range of k?