Representation Theory

Ricky Roy
Math 434
University of Puget Sound

May 2, 2010

Introduction

In our study of group theory, we set out to classify all distinct groups of a given order up to isomorphism. In doing so, we were able to develop tools and language which allowed us to say that two seemingly different groups were actually the "same," in that they have the same structure. This information was all contained in the group Cayley table, a construct which will have an important counterpart in representation theory, called the character table.

Informally, a representation of a group is a "copy" of the group structure within $GL(m)$, the group of $m \times m$ invertible matrices. These matrix representations must satisfy the group Cayley table under matrix multiplication. An initial characteristic to note about representations then is that they are basis dependent since they are matrices. The desire for uniqueness of representations of a group allows us to formulate the fundamental problem of representation theory: classify all representations of a group G up to isomorphism. This paper focuses only on the general linear group over the complex numbers, $GL(m;\mathbb{C})$, because this alone is a sufficiently difficult task. Furthermore, we will be focusing on representations of finite groups.

The goal of this paper is to develop and prove most of the crucial and fundamental results of representation theory, giving the reader sufficient tools to understand and generate the most important construct of representation theory, the character table.

Representation Theory of Finite Groups

Preliminaries

Definition 1. A **representation** (ρ, V) of a group G on a finite-dimensional complex vector space V is a homomorphism $\rho : G \mapsto GL(V)$.

Note that the requirement that a representation preserve the group table structure is contained in the fact that a representation is a homomorphism. For any $g_1, g_2 \in G$, where $g_1g_2 = g_3 \in G$,

$$\begin{aligned}\rho(g_1)\rho(g_2) &= \rho(g_1g_2)\\ &= \rho(g_3)\end{aligned}$$

The first line follows from the fact that ρ is already a homomorphism. A representation is called faithful if each group element is represented by a distinct matrix. More formally, a representation is **faithful** if ρ is one-to-one.

Example 1. Given any group G, we can construct the **trivial representation** by sending every group element to the complex scalar 1. The group Cayley table is then satisfied, but the representation is "as far from faithful as possible."

The power behind choosing a vector space V over which to define the representation is that we can use all of the already developed machinery for vector spaces to characterize representations. As noted in the introduction, the matrix representations will be basis dependent, so we will want to know when we are dealing with "isomorphic representations."

Definition 2. Two representations of a group G, (ρ_1, V_1) and (ρ_2, V_2), are said to be **isomorphic representations** if there exists an isomorphism of vector spaces $\phi : V_1 \mapsto V_2$ which commutes with the action of G on V, meaning:

$$\phi(\rho_1(g)\mathbf{v}) = \rho_2(g)(\phi(\mathbf{v})) \quad \forall g \in G, \mathbf{v} \in V.$$

If ϕ is not invertible (in which case we do not have an isomorphism of representations), it is referred to as an **intertwining operator** or **G-linear map**. In the case where it is invertible, rearranging the above expression gives the explicit form for changing bases between isomorphic representations,

$$\rho_2 = \phi \circ \rho_1 \circ \phi^{-1}.$$

The above line can also be viewed as a way of constructing a seemingly different representation given one which we already have. This will amount to a similarity transformation of matrices.

One of the biggest topics in representation theory is reducibility of a representation. In other words, given a representation, could we encode the same information about the group with less dimensions, as a direct sum of smaller dimensional vector spaces?

Definition 3. A **subrepresentation** of (ρ, V) is an invariant subspace $W \subseteq V$ under the action of G (also called a G-invariant subspace), which says that for all $\mathbf{w} \in W$ and $g \in G$, $\rho(g)\mathbf{w} \in W$.

A representation is then called **irreducible** if it contains no proper invariant subspaces (subrepresentations), and is defined to be **completely reducible** if it can be expressed as a direct sum of irreducible representations. In terms of the matrix representations themselves, reducibility amounts to simultaneous block diagonalization of $\rho(g)$ for all $g \in G$, via the similarity transformation discussed above.

The Complete Reducibility Theorem

Definition 4. The **direct sum** of two representations (ρ_1, V_1) and (ρ_2, V_2) is expressed as $(\rho_1 \oplus \rho_2, V_1 \oplus V_2)$, where $\rho_1 \oplus \rho_2$ has block diagonal action on $V_1 \oplus V_2$.

The following theorem allows us to establish the *Complete Reducibility Theorem*, one of the major theorems of representation theory. We will then be able to take any representation of a group, and express it as a direct sum of irreducible representations.

Theorem 1. If W is a subrepresentation of (ρ, V), then there exists a complementary invariant subspace W' such that $W \oplus W' = V$.

Proof. We will take two facts from linear algebra and prove the part that is most important for the *complete reducibility theorem.* First, given a subspace W of V, there is a complement W' such that the direct sum of the two equals V. Second, given a representation ρ', we can construct an isomporphic representation ρ which lives in $U(V)$, the group of unitary transformations of V.

The important property for us to prove is that this complement is G-invariant. To do so, we must show that given $\mathbf{w}' \in W'$, $\rho(g)\mathbf{w}' \in W'$ for all $g \in G$. Using that W and W' are complements, we know that for $\mathbf{w} \in W$, $\langle \mathbf{w}' | \mathbf{w} \rangle = 0$. Because $\rho(g)$ is unitary, this says $\langle \rho(g)\mathbf{w}' | \rho(g)\mathbf{w} \rangle = 0$ for any $g \in G$. Using the invariance of W, we can say $\rho(g)\mathbf{w} = \mathbf{w}_1 \in W$, for arbitrary $g \in G$. Then since, $\langle \rho(g)\mathbf{w}' | \mathbf{w}_1 \rangle = 0$, $\rho(g)\mathbf{w}'$ must live in W', so W' is a G-invariant subspace. □

We can now easily see how repeated decomposition of a given representation into direct sums of G-invariant supspaces will eventually end in expressing that representation as a direct sum of irreducible subrepresentations.

Corollary 1. The Complete Reducibility Theorem: Any representation (ρ, V) can be decomposed into a direct sum of irreducible representations.

Schur's Lemma

The second major theorem we will discuss is known as Schur's Lemma, which establishes uniqueness in terms of isomorphism classes for irreducible representations.

Schur's Lemma 1. If (ρ_1, V_1) and (ρ_2, V_2) are irreducible representations of a group G, then any nonzero homomorphism $\phi : V_1 \mapsto V_2$ is an isomorphism.

Proof. Assuming ϕ is nonzero, we can write $\mathbf{v}_2 = \phi(\mathbf{v}_1) \in V_2$ for some $\mathbf{v}_1 \in V_1$. We can then say that $\rho_2(g)(\mathbf{v}_2) = \rho_2(g)(\phi(\mathbf{v}_1))$, which by the intertwining property of maps between representations, gives that $\rho_2(g)(\mathbf{v}_2) = \phi(\rho_1(g)(\mathbf{v}_1)) \in \phi(V_1)$. This tells us that $\mathrm{Im}(\phi) = \phi(V_1)$ is a G-invariant subspace of V_2, which by the irreducibility says that $\phi(V_1) = V_2$. This gives us onto. To attack one-to-one, we will show that the kernel is trivial.

For $\mathbf{v}_1 \in Ker(\phi)$, $\phi(\rho_1(g)(\mathbf{v}_1)) = \rho_2(g)(\phi(\mathbf{v}_1)) = \rho_2(g)(0) = 0$ for all $g \in G$. Therefore, the kernel of ϕ is a G-invariant subspace of V_1, and since the range of ϕ is not empty, the kernel must be a proper subspace. By irreducibility of ρ_1, $Ker(\phi) = 0$, and ϕ is an isomorphism. □

Schur's Lemma tells us that given two irreducible representations, we can either find that they are exactly the same (isomorphic), or that the only map between them is the zero map. We can now refer to **isomorphism classes of irreducible representations**. The following theorem is often contained in a larger single version of Schur's Lemma, hence the redundancy of title.

Schur's Lemma 2. Let (ρ, V) be an irreducible representation of a group G and $\phi : V \mapsto V$ a nonzero homomorphism. Then $\phi = \lambda \cdot \mathrm{Id}$ for some $\lambda \in \mathbb{C}$, where Id is the identity transformation.

Proof. By Schur's Lemma 1, ϕ is an isomorphism. Let λ be an eigenvalue of ϕ. Take $\phi - \lambda \cdot \mathrm{Id}$, which is a homomorphism from V to V with zero determinant (set the characteristic polynomial equal to zero to find λ). A map with a zero determinant cannot be an isomorphism. Therefore, by Schur's Lemma 1 $\phi - \lambda \cdot \mathrm{Id} = 0$, and $\phi = \lambda \cdot \mathrm{Id}$ for some $\lambda \in \mathbb{C}$. Note the necessity for the representation to be over an algebraically closed field so that the eigenvalues are always computable. □

We now have enough machinery to characterize the representations of all abelian groups.

Corollary 2. Any irreducible complex representation of an abelian group is 1-dimensional.

Proof. Let (ρ, V) be an irreducible complex representation of G. Since G is abelian, we know that $\rho(g)\rho(h)\mathbf{v} = \rho(gh)\mathbf{v} = \rho(hg)\mathbf{v} = \rho(h)\rho(g)\mathbf{v}$ for all $\mathbf{v} \in V$. By Schur's Lemma 2, $\rho(g)\mathbf{v} = c\mathbf{v}$ for any $g \in G$, where c is some complex scalar. Therefore, every subspace of V will be G-invariant, or in other words is a subrepresentation. Irreducibility of V implies that the only subrepresentations are the trivial spaces $\{0\}$ and V itself. Any assertion that $\dim V > 1$ requires V to have a non-trivial subspace (for example, the span of a proper nonzero subset of a basis for V is a non-trivial subspace of V), which contradicts the previous statement. Therefore, V is 1-dimensional. $\square$

Corollary 3. Any irreducible complex representation of a cyclic group is 1-dimensional.

Proof. Cyclic groups are abelian. By the preceding corollary, we are done. $\square$

Example 2. Find an irreducible representation of the cyclic group G of order 7 (isomorphic to $\mathbb{Z}_7$ under addition). By Corollary 3, any irreducible representation is 1-dimensional. Furthermore, we know that $\rho(1)^7 = \rho(1^7) = \rho(1) = 1$. Therefore, if x is a complex scalar representation of 1_G, it must be a root of the polynomial $x^7 - 1$. Therefore, it must be a 7^{th} root of unity:

$$\rho(1) = e^{2\pi i n/7}$$

for any given $n \in \mathbb{Z}_7$. The rest of the group representation is then generated by taking the seven powers of $\rho(1)$. If $n = 0$, then $\rho(1) = 1$ and the representation is the trivial representation. Otherwise, the representation is faithful.

Characters

Now that we have established a distinction between irreducible representations of a given group, we will want to be able to find a property of the distinct irreducibles which will be the same for two isomorphism representations. Recall from our discussion of the intertwining operator that if we have two isomorphic representations, we can perform a simultaneous similarity transformation on one representation to get to the other. The invariance of the trace under a similarity transformation (elementary linear algebra result) makes it a strong candidate for this distinctive property.

Definition 5. The **character** of a representation ρ is a function χ_ρ on G defined by $\chi_\rho(g) = \mathrm{Tr}(\rho(g))$.

Theorem 2. If χ_ρ is the character of a representation ρ of a group G, then χ_ρ has the following properties:

1. The character is the same for elements related by conjugation. $\chi_\rho(g) = \chi_\rho(hgh^{-1})$ for all $g \in G$ and $h \in G$.

2. $\chi_\rho(1) = \dim V$.

3. $\chi_\rho(g^{-1}) = \chi_\rho(g)^*$, where * denotes the complex conjugate (or conjugate transpose for a matrix).

4. For two representations V and W, $\chi_{V\oplus W} = \chi_V + \chi_W$.

Proof. (1) follows from the fact that $\rho(hgh^{-1}) = \rho(h)\rho(g)\rho(h^{-1})$. Then $\mathrm{Tr}(\rho(h)\rho(g)\rho(h^{-1})) = \mathrm{Tr}(\rho(hh^{-1})\rho(g)) = \mathrm{Tr}(\rho(g))$. For (2), $\rho(1) = \mathrm{Id}_V$. For (3), remember we can find a unitary representation ρ' given ρ, then $\rho'(g^{-1}) = \rho'(g)^{-1} = \rho'(g)^*$. Lastly for (4), direct product of representations corresponds to block diagonal matrices. Therefore the trace of the resulting matrix will be the sum of the two original traces. $\square$

Definition 6. The **inner product** of two complex functions ψ and ϕ on G is defined to be:

$$\begin{aligned} \langle\psi|\phi\rangle &= \frac{1}{|G|}\sum_{g\in G}\psi(g)^*\phi(g) \\ &= \frac{1}{|G|}\sum_{C}|C|\psi(C)^*\phi(C), \end{aligned}$$

where the second line denotes summing over conjugacy classes and weighting each by the number of elements in it.

The following theorem, called the *Orthogonality of Characters* is one of the most important results in representation theory, and helps us reach virtually every result from here on out. The proof is too lengthy for this paper, and would add little to our purpose, but can be found in any of the sources cited in the bibliography.

Theorem 3. Orthogonality of Characters:

1. If V is an irreducible representation, then $\langle\chi_V|\chi_V\rangle = 1$.

2. If V and W are irreducible representations of G and not isomorphic, $\langle\chi_V|\chi_W\rangle = 0$.

In our Abstract Algebra course, we have come across a specific representation of utmost importance in representation theory. This is called the **regular representation**, which we encountered as a result of Cayley's theorem. The character of the regular representation will lead us to a couple important results. From Cayley's theorem, we found that any group is isomorphic to a subgroup of the permutation group on $|G|$ symbols. The regular representation of an element $g \in G$ is then the permutation corresponding to the function $\lambda_g(a) = ga$, which takes $a \in G$ to another element in G. This can be encoded as a matrix in $GL(\mathbb{C}^{|G|})$ of 1's and 0's by assigning basis vectors according to the order of group elements which g acted upon. An illustrative example will help to have a clear picture of the regular representation.

Example 3. Find the regular representation ρ of $J \in Q$, where Q is the quaternions. Take the ordered set $Q = \{1, -1, I, -1, J, -J, K, -K\}$ and label each element by its index, ranging from 1 to 8. Then,

$$\lambda_J = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 \\ 5 & 6 & 8 & 7 & 2 & 1 & 3 & 4 \end{pmatrix}$$

Then, in the basis where the i^{th} group element corresponds to $\mathbf{e}_i$, the i^{th} standard basis vector, we have:

$$\rho(J) = \begin{pmatrix} 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{pmatrix}$$

Since a 1 along the diagonal corresponds to an element being fixed, and we know the identity element to be unique ($ga = a \Rightarrow g = 1$), we have proven the following theorem.

Theorem 4. The character of the regular representation is:

$$\chi_{\text{reg}}(g) = \begin{cases} |G| & \text{if } g = 1 \\ 0 & \text{otherwise} \end{cases}$$

Theorem 5. The number of times an irreducible representation V appears in the decomposition of W, called the **multiplicity**, is $\langle \chi_V | \chi_W \rangle$.

Proof. This proof relies on the orthogonality theorem. Recall from Theorem 2 (4) that the character of a representation which is the direct sum of irreducible representations is the sum of the characters of the individual representations. By the *Complete Reducibility Theorem*, we can decompose W into a direct sum of n irreducible representations, and therefore χ_W as a sum of the individual characters of the corresponding n irreducible representations. So we have,

$$\begin{aligned} \langle \chi_V | \chi_W \rangle &= \langle \chi_V | \chi_{W_1} + \chi_{W_2} + \dots + \chi_{W_n} \rangle \\ &= \langle \chi_V | \chi_{W_1} \rangle + \langle \chi_V | \chi_{W_2} \rangle + \dots + \langle \chi_V | \chi_{W_n} \rangle \qquad \text{(linearity from Definition 6)} \end{aligned}$$

From the orthogonality theorem,

$$\langle \chi_V | \chi_{W_i} \rangle = \begin{cases} 1 & \text{if } W_i = V \\ 0 & \text{if } W_i \neq V \end{cases}$$

So we have proved that if V shows up k times in the decomposition of W, then $\langle \chi_V | \chi_W \rangle = k$. □

The next two theorems should justify my assertion of the importance of the regular representation.

Theorem 6. The multiplicity of any irreducible representation (ρ, V) of G in the regular representation is equal to the dimension of V.

Proof.

$$\begin{aligned}
\langle \chi_V | \chi_{\text{reg}} \rangle &= \frac{1}{|G|} \sum_g \chi_V(g)^* \chi_{\text{reg}}(g) \\
&= \frac{1}{|G|} \chi_V(1)^* \chi_{\text{reg}}(1) && \text{(Theorem 4)} \\
&= \frac{|G|}{|G|} \chi_V(1)^* && \text{(Theorem 4)} \\
&= \chi_V(1) && \text{(Theorem 2 part 3)} \\
&= \text{Tr}(I_V) \\
&= \dim V
\end{aligned}$$

□

Lemma 1. $\chi_{\text{reg}} = \sum_i \dim V_i \cdot \chi_{V_i}$, where i ranges over the isomorphism classes of irreducible representations.

Proof. We know that the regular representation can be decomposed into a direct sum of all the distinct irreducible representations of the group, V_i, and that their multiplicity is equal to the number of times they appear in the direct sum. Therefore we can write the character of the regular representation as a sum of the irreducible characters, weighted by their multiplicity (or dimension by the previous theorem):

$$\chi_{\text{reg}} = \sum_i \dim V_i \cdot \chi_{V_i}$$

□

Theorem 7. The sum of the squares of the dimensions of all distinct irreducible representations of G is equal to the order of the group:

$$\sum_i \dim V_i^2 = |G|.$$

Proof.

$$\begin{aligned}
|G| &= \frac{|G| \cdot |G|}{|G|} \\
&= \frac{1}{|G|} \chi_{\text{reg}}(1)^* \chi_{\text{reg}}(1) && \text{(Theorem 4)} \\
&= \frac{1}{|G|} \sum_g \chi_{\text{reg}}(g)^* \chi_{\text{reg}}(g) && \text{(Theorem 4)} \\
&= \langle \chi_{\text{reg}} | \chi_{\text{reg}} \rangle && \text{(Definition 6)} \\
&= \sum_i \dim^2 V_i \, \langle \chi_{V_i} | \chi_{V_i} \rangle && \text{(Lemma 1)} \\
&= \sum_i \dim^2 V_i && \text{(Orthogonality Theorem)}
\end{aligned}$$

□

At this point, we are equipped with almost every piece of information necessary to classify all representations of a given group. The piece missing is the number of distinct irreducible representations of a given group. The corresponding theorem is called *The Completeness of Irreducible Characters*, and states that the number of irreducible representations is equal to the number of conjugacy classes in the group. The proof of this theorem is beyond the scope of this paper (can be found in any of the sources listed in the bibliography), so we will stop at simply proving that the number of irreducible representations is less than or equal the number of conjugacy classes.

Theorem 8. The number of irreducible representations of a group G is less than or equal to the number of conjugacy classes in G.

Proof. Denote the number of conjugacy classes of G by n. Utilizing the *kroncker delta* symbol δ_{ij}, which equals 1 if $i = j$ and 0 otherwise, we can say:

$$\begin{aligned}\delta_{ij} &= \langle \chi_i | \chi_j \rangle \\ &= \frac{1}{|G|} \sum_{k=1}^{n} |C_k| \chi_i(C_k)^* \chi_j(C_k)\end{aligned}$$

The χ_i's can be treated as n-dimensional vectors with its k^{th} component corresponding to the character of the k^{th} class. In this guise, the equality above is a relation of orthonormality between these vectors. Herein lies the upper limit on the number of irreducible representations. Since this is an n-dimensional vector space, there must be less than or equal to n vectors in an orthonormal set. Therefore, the number of values which i or j can take on, which corresponds to the total number of irreducible representations, is less than or equal to n. □

The Character Table

We now have all the information necessary to characterize all irreducible representations of a finite group G. We will list this information here for ease of reference:

- The character is the same for all elements in a given conjugacy class.
- $\sum_i \dim^2 V_i = |G|$, where i runs over distinct irreducible representations.
- $\left\langle \chi_{V_i} | \chi_{V_j} \right\rangle = \delta_{ij}$, where i and j both run over distinct irreducible representations.
- The number of irreducible representations equals the number of conjugacy classes.

Definition 7. Given a group G with n conjugacy classes, its **character table** is the $n \times n$ matrix whose ij^{th} entry is the character of the j^{th} conjugacy class in the i^{th} irreducible representation.

Informally, we will draw up character tables by labeling the rows with the "name" of the irreducible representation, and the column with the conjugacy class. The last two examples are meant to summarize and connect the main points covered in this paper on representation theory.

Example 4. Find the character table for the cyclic group $|G|$ of order 7. For abelian (and therefore cyclic) groups, conjugacy classes consist of a single element because $hgh^{-1} = hh^{-1}g = g$ for all $g, h \in G$. Therefore, there must be as many irreducible representations as there are group elements (i.e. the character table is a 7×7 matrix). Furthermore, since all irreducible representations of abelian groups are 1-dimensional, the characters (or traces) are just the scalar representations themselves. As we found in Example 2, the 7 possible representations of 1_G are given by $e^{2\pi in/7}$ for $n \in \mathbb{Z}_7$. For ease of ensuing notation, let's call $\omega = e^{2\pi i/7}$. For each of the irreducible representations Γ^l, the l will correspond to the choice of $\Gamma^l(1) = \omega^l$. Remembering that $\Gamma^l(k) = \Gamma^l(1^k) = \Gamma^l(1)^k = \omega^{lk} = \omega^{lk \bmod 7}$, we get the following character table:

	0	1	2	3	4	5	6
E	1	1	1	1	1	1	1
Γ^1	1	ω^1	ω^2	ω^3	ω^4	ω^5	ω^6
Γ^2	1	ω^2	ω^4	ω^6	ω^1	ω^3	ω^5
Γ^3	1	ω^3	ω^6	ω^2	ω^5	ω^1	ω^4
Γ^4	1	ω^4	ω^1	ω^5	ω^2	ω^6	ω^3
Γ^5	1	ω^5	ω^3	ω^1	ω^6	ω^4	ω^2
Γ^6	1	ω^6	ω^5	ω^4	ω^3	ω^2	ω^1

For the second example, we will do the less trivial case of the Quaternions.

Example 5. Find the character table for the Quaternions. It is often the case with small non-abelian groups to use ad hoc methods of finding the irreducible representations, which is what will be illustrated here. First, we know that the total number of irreducible representations equals the number of conjugacy classes. In the case of the Quaternions, the set of conjugacy classes is $\{\{1\}, \{-1\}, \{I, -I\}, \{J, -J\}, \{K, -K\}\}$. Therefore, we know we will have 5 irreducible representations. Furthermore, we know that the sum of the squares of the dimensions must equal the order of the group. So we need 5 integers whose squares sum to 8. The only way to accomplish this is $2^2 + 1^2 + 1^2 + 1^2 + 1^2 = 8$. This means we will have 1 irreducible representation of degree 2 and 4 irreducible representations of degree 1.

To tackle the 1-dimensional representations, consider the subgroups of the Quaternions; $\{1\}, \{\pm 1\}, \{\pm 1, \pm I\}, \{\pm 1, \pm J\}, \{\pm 1, \pm K\}$, and the whole group. Since a representation is a homomorphism, its kernel must be a subgroup of the whole group. If the kernel is the entire group, we have the identity transformation, E. Next, we can choose any of the maximal normal subgroups to be the kernel. If we choose the kernel to be $\{\pm 1, \pm I\}$, then we define ρ_I to act such that $\rho_I(1) = \rho_I(-1) = \rho)_I(I) = \rho_I(-I) = 1$. For this irreducible representation to not be the same as the trivial representation, it must send $\pm J$ and $\pm K$ to something other than 1, but we also must satisfy $\rho_I(\pm J)^2 = \rho_I(J^2) = \rho_I(-1) = 1$, and the same for $\pm K$. Therefore, our only choice is to say that $\rho_I(J) = \rho_I(-J) = \rho)_I(K) = \rho_I(-K) = -1$. Repeating this process for the kernel equaling the other two maximal normal subgroups gets the remaining two 1-dimensional representations.

The 2-dimensional irreducible representation is a bit more tricky. Let's try and find a faithful representation, ρ_{2-d}. If we do find a faithful representation, we know it must be irreducible since none of the 1-dimensional representations were faithful. We know that

$\rho_{2-d}(1) = I_{2\times 2}$. Then let's specify $\rho_{2-d}(-1) = -I_{2\times 2}$. Now let's take a "leap" and say (as Wolfgang Pauli would):

$$\rho_{2-d}(I) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad \rho_{2-d}(J) = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \qquad \rho_{2-d}(K) = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

By our specification of $\rho_{2-d}(-1)$ then, the $-I, -J$, and $-K$ representations will be the negatives of the above matrices. Check that they satisfy the group structure (they do). Therefore, we have found an irreducible faithful representation of the Quaternions of dimension 2, and we get the following character table:

	$\{1\}$	$\{-1\}$	$\{I, -I\}$	$\{J, -J\}$	$\{K, -K\}$
E	1	1	1	1	1
ρ_I	1	1	1	-1	-1
ρ_J	1	1	-1	1	-1
ρ_K	1	1	-1	-1	1
ρ_{2-d}	2	-2	0	0	0

It is reassuring that the rows and columns of the character table matrices in the previous two examples are orthonormal sets of vectors as we would expect. The first example is a little more difficult and requires a bit of complex analysis, but the second is clear as day.

Conclusion

Let us restate the initial goals of this paper and what we have accomplished. We set out to motivate and develop the major ideas behind basic representation theory of finite groups, the main problem of is to classify all irreducible representations of a given group by the group's character table. We did just that for *all* finite abelian and cyclic groups, and did an illustrative example with an interesting non-abelian group of order 8, the Quaternions. We have covered just a small subset of general representation theory, but I believe an enlightening subset at that.

Bibliography

[1] W. Fulton, J. Harris, *Representation Theory: A First Course*, Springer-Verlag New York Inc., New York, NY, 1991.

[2] M. Burrow, *Representation Theory of Finite Groups*, Dover Publications, Inc., New York, 1993.

[3] C. Teleman, *Representation Theory*, 2005, available at http://math.berkeley.edu/∼teleman/math/RepThry.pdf.

[4] P. Etingof, O. Golberg, S. Hensel, T. Liu, A. Schwendner, E. Udovina and D. Vaintrob, *Introduction to Representation Theory*, 2009, available at http://www-math.mit.edu/∼etingof/replect.pdf.