Darek's notes on representation theory

(These notes should copy what I said in the lectures. They have been written hastily, so they may contain typos etc. But the math really should be OK.)

1 Modules

In its heart, representation theory is concerned with studying how groups can act on vector spaces. For that reason, the first structure we should introduce is a vector space that also "knows what to do" if elements of a group act on it. Of course, the group should respect the structure of the vector space too, so if $\boldsymbol{v}$ and $\boldsymbol{w}$ are two vectors in such a vector space and g is an element of the group G, we should have $g(\boldsymbol{v}+\boldsymbol{w}) = g\boldsymbol{v}+g\boldsymbol{w}$. Similarly, if α is a constant, we should have $g(\alpha\boldsymbol{v}) = \alpha g\boldsymbol{v}$.[1] Such a vector space is called a *G-module*.

The formal definition is as follows: a vector space V is called a *G-module* if for every $g \in G$ and $\boldsymbol{v} \in V$, a unique $g\boldsymbol{v} \in V$ is defined in such a way that: first, $g(\alpha\boldsymbol{v} + \beta\boldsymbol{w}) = \alpha g\boldsymbol{v} + \beta g\boldsymbol{w}$ (for all $v, w \in V$ and $\alpha, \beta \in \mathbb{C}$); second, if $g = g_1g_2$, then $g\boldsymbol{v} = g_1(g_2\boldsymbol{v})$ (for all $g_1, g_2 \in G$); third, $e\boldsymbol{v} = \boldsymbol{v}$ (where e is the identity element of the group).

As you can see, it's precise, but the essence is that "a *G-module* is a vector space on which elements of group G act as we would intuitively expect".

Here's an example: consider the group S_3 of permutations of three symbols. It has six elements: the identity permutation (), the three two-cycles (12), (23) and (31), and the two three-cycles (123) and (132). We take the vector space $\mathbb{R}^3$ with an orthonormal basis $\boldsymbol{e}_1$, $\boldsymbol{e}_2$ and $\boldsymbol{e}_3$ and we say that each permutation of S_3 acts on this space by permuting the basis vectors accordingly. So (12) will switch $\boldsymbol{e}_1$ and $\boldsymbol{e}_2$ (geometrically, this is a reflection through some plane), (123) will permute the three vectors cyclically (which corresponds to the rotation by $2\pi/3$ around the axis $\boldsymbol{e}_1 + \boldsymbol{e}_2 + \boldsymbol{e}_3$) etc. This is an S_3-module.

This module is a bit peculiar. It is easy to see that the vector $\boldsymbol{e}_1 + \boldsymbol{e}_2 + \boldsymbol{e}_3$, or any multiple of it, is unaffected by any of the operations in S_3. This means that the one-dimensional subspace spanned by this vector is also an S_3-module in itself — its vectors still have the ability to be acted upon by elements of S_3. Moreover the result of any such operation will belong to the subspace, and it can never "leave" it. Hence this subspace is actually a *submodule* of the original module.

2 (Ir)reducibility of modules

Obviously, any module has two trivial submodules: the zero module (that contains only the zero vector and nothing else), and the module itself. These aren't very interesting. If a module as only these two submodules, it is called *irreducible*. On the other hand, if a module has a non-trivial submodule, we say it is *reducible*.

In our example of an S_3-module, we actually have a bit more: not only the subspace spanned by $\boldsymbol{e}_1+\boldsymbol{e}_2+\boldsymbol{e}_3$ is a submodule, but the "rest" of the vector space (a two-dimensional subspace spanned by $\boldsymbol{e}_1-\boldsymbol{e}_2$ and $\boldsymbol{e}_2-\boldsymbol{e}_3$) is a submodule as well. Hence we can write the module as a *direct sum*[2] of two smaller, but still non-trivial submodules. In this case, we say the module is *completely reducible*.

It may seem that this is just playing with words and that "reducibility" and "complete reducibility" must be the same thing. And often, it is. However, consider this example: we take the vector space $\mathbb{R}^2$, pick some basis in it, and we act on it with the group $\mathbb{Z}^3$ as follows: each integer k acts as a matrix $\begin{pmatrix} 1 & k \\ 0 & 1 \end{pmatrix}$. Since $\begin{pmatrix} 1 & k \\ 0 & 1 \end{pmatrix}\begin{pmatrix} 1 & n \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & k+n \\ 0 & 1 \end{pmatrix}$, this vector space is a $\mathbb{Z}$-module. It is obviously reducible: the basis vector $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ is unchanged by any of those matrices, so the subspace generated by it is a submodule. However, the

[1] This will probably feel more natural if we realize that the elements of groups often correspond to some geometrical operations like rotations, reflections or translations.

[2] A sum of vector spaces $V_1 + V_2$ is just a vector space formed by taking all possible linear combinations of vectors, one from V_1 and one from V_2. If the sum is *direct*, it just means that V_1 and V_2 are disjoint — they have just the zero vector in common.

[3] The group of all integers with addition.

"rest" of the vector space (the subspace generated by $\begin{pmatrix}0\\1\end{pmatrix}$) is not a submodule, because each group element makes these vectors "leak" into the $\begin{pmatrix}1\\0\end{pmatrix}$ direction. Hence this module is **NOT** completely reducible!

By the way: by now, you probably feel that this (ir)reducibility of modules must have something to do with the (ir)reducibility of representations. And it does. We use the modules because they offer a much clearer picture. Thanks to them, we see that the (ir)reducibility is the property of the module, i. e. of how the group acts on a vector space, and not of the representation matrices. The problem with matrices lies in the fact that we must always choose a basis in order to construct them; otherwise we wouldn't know which numbers to fill in. But that means that we are introducing extraneous information that obscures the problem and that we must actively grapple with. This only leads to unnatural definitions like "a bunch of matrices is completely reducible if we can change basis in such a way that they all simultaneously go to the block-diagonal form". Hence it is (in my opinion) better to use the language of modules.

3 Group algebra

Before we continue, we make a slight generalization. In a module, we can write things like $g_1\boldsymbol{v} + 2g_2\boldsymbol{v}$. This could be also understood as a formal sum of group elements, $g_1 + 2g_2$, acting on the vector $\boldsymbol{v}$. Of course, someone could object that groups contain no notion of addition and scalar multiplication, and so these sums do not make any sense. That is true to some extent.[4] However, once we have some vector space on which the group elements can act, we can easily act on it with such sums as well. In fact, we use sums of this type pretty often in physics — for instance, we often write things like $(\frac{\partial}{\partial x} - 3\frac{\partial}{\partial y})f$ and understand them to be the same as $\frac{\partial f}{\partial x} - 3\frac{\partial f}{\partial y}$. Here we use the same principle.

These formal sums can also be multiplied in an intuitive fashion. For instance, $(g_1 - g_2)(2g_1^{-1} + g_3) = 2 + g_1g_3 - 2g_2g_1^{-1} - g_2g_3$ — we assume the distributive law, numbers go to the front and the group elements multiply according to their multiplication table. All possible formal linear combinations of group elements, $\sum c_n g_n$, with the intuitive multiplication rules, make up an object called the *group algebra* of that group. As physicists, we assume that the coefficients of the linear combination can be any complex numbers, and we denote the group algebra of G simply by $\mathbb{C}G$.

Now we can also slightly extend our idea of a G-module: we can allow acting not only with the group elements, but with the formal linear combinations from $\mathbb{C}G$ as well. If we once more demand that these linear combinations do the intuitive thing, i. e. $(\sum c_n g_n)\boldsymbol{v} = \sum c_n(g_n\boldsymbol{v})$, we upgrade our G-module to a "$\mathbb{C}G$-module". Since this is so easy and natural to do, we will work with these "algebra-modules" from now on. It may seem like a useless formal trick, but it will come in handy later.

4 Maschke's theorem

Let's go back to the reducibility of modules now. We showed that reducibility and complete reducibility were different things. However, we can immediately forget about it, because there is an important result called *Maschke's theorem* which simply says that *for any finite group G, any $\mathbb{C}G$-module is always completely reducible*, so we always have the nice case.

What's more, the idea that makes it work is quite simple, and it can be gleaned from the S_3 example that we started with. That particular S_3-module decomposes to a direct sum of the one-dimensional submodule (a line in $\mathbb{R}^3$) and its orthogonal complement (the plane perpendicular to it). It turned out that the group action was unable to violate the orthogonality.

Hence we consider a $\mathbb{C}G$-module with a scalar product[5] $\langle\cdot,\cdot\rangle$, and we use it to build an *invariant scalar product* $\{\cdot,\cdot\}$ that does not change if we act with the same group element on both its operands, i. e. $\{g\boldsymbol{v}, g\boldsymbol{w}\} = \{\boldsymbol{v}, \boldsymbol{w}\}$ for any $g \in G$. If we can manage to do it, we win. If we find a submodule, we can use this

[4] This is why I say they are "formal" sums — they are objects obtained by writing some group elements with some numbers in front of them and sticking plus signs in between them, nothing more.

[5] Remember, the modules are just vector spaces, so they can have scalar product. If the module doesn't have it, we can always make one up by picking any basis and defining the scalar product by postulating that the basis vectors satisfy $\langle \boldsymbol{e}_k, \boldsymbol{e}_\ell\rangle = \delta_{k\ell}$.

invariant product to get its orthogonal complement. Moreover, if we take any vector $\boldsymbol{v}$ from the submodule and any $\boldsymbol{w}$ from the orthogonal complement, they will stay orthogonal no matter what group element is applied to them (since the product is invariant), and the vectors in the complem ent cannot leak into our submodule.

So we just have to find an invariant scalar product. We do it as follows:

$$\{\boldsymbol{v}, \boldsymbol{w}\} := \frac{1}{|G|} \sum_{g \in G} \langle g\boldsymbol{v}, g\boldsymbol{w} \rangle.$$

With this definition, $\{h\boldsymbol{v}, h\boldsymbol{w}\} = \frac{1}{|G|} \sum \langle gh\boldsymbol{v}, gh\boldsymbol{w} \rangle$. But for any group, each row and column of the multiplication table contains each group element exactly once. This means that if g goes through the whole group, gh does too, and the sum is unchanged, so we truly have $\{\boldsymbol{v}, \boldsymbol{w}\} = \{h\boldsymbol{v}, h\boldsymbol{w}\}$ for any $h \in G$. The product that we defined really is invariant, and Maschke's theorem is proved.

This is a result that will make our job much easier, because it can be applied repeatedly even to submodules of our module. Each time, we find out that reducible submodules (or their reducible submodules etc.) decompose into direct sums. This means that in the end, every $\mathbb{C}G$-module must decompose to a direct sum of irreducible submodules that do not interact with each other at all when group elements are applied. It turns out that for each group G, there is only limited number of different types of irreducible modules, and every $\mathbb{C}G$-module must be just a direct sum of some of these.

5 Homomorphisms

Now that we see that we are getting nice results for $\mathbb{C}G$-modules, we ask the question that mathematicians always ask when they encounter a new structure: Suppose that we have different objects equipped with the structure. What kind of maps can we make between them so that the structure is not ruined? Such maps always get the traditional name *homomorphism* (and we already saw one example of it, since we proved some properties of homomorphisms between groups).

Now suppose we have two $\mathbb{C}G$-modules, U and V, and some map $\varphi : U \to V$ between them. What properties should φ have in order to deserve the name of homomorphism between the $\mathbb{C}G$-modules? It must be a map that preserves their nice properties. So we must recollect what these properties are: firstly, the modules are *vector spaces*, secondly, they *know how* $\mathbb{C}G$ *acts on them* (and the elements of $\mathbb{C}G$ act in a certain consistent and intuitive manner).

A homomorphism between $\mathbb{C}G$-modules should preserve both of these things, so we define it as a *linear map* (i. e. it preserves the vector space structure) that also satisfies $\varphi(g\boldsymbol{u}) = g\varphi(\boldsymbol{u})$ (for any $\boldsymbol{u} \in U$ and any $g \in \mathbb{C}G$; of course, the action of g on $\varphi(\boldsymbol{u})$ on the right-hand side is that of the "destination" module V).[6]

Maps that satisfy the two conditions above (linear map + $\varphi(g\boldsymbol{u}) = g\varphi(\boldsymbol{u})$) should be called "$\mathbb{C}G$-module homomorphisms". But since that is quite a mouthful, and we will not consider homomorphisms of other objects than modules, we just shorten the name to a "$\mathbb{C}G$-homomorphism".

Now that we know what the $\mathbb{C}G$-homomorphisms are, we can start studying their properties. We know from linear algebra that each linear map defines two important subspaces of the "source" and "result" modules. It is the *kernel*, $\operatorname{Ker}\varphi$, which is the subspace of everything in U that gets sent to the zero in V, and the *image*, $\operatorname{Im}\varphi$, which is the subspace of all possible results that φ can give. We should have a closer look on what the group action does to these subspaces.

Let's have look at the kernel first. Suppose that some vector $\boldsymbol{k} \in U$ belongs to the kernel, which means that $\varphi(\boldsymbol{k}) = \boldsymbol{0}$. Now what happens if we act on $\boldsymbol{k}$ by some element $g \in \mathbb{C}G$? The map φ will send it to $\varphi(g\boldsymbol{k})$, but, since φ is a homomorphism, this must be the same thing as $g\varphi(\boldsymbol{k}) = g\boldsymbol{0} = \boldsymbol{0}$. So if $\boldsymbol{k}$ is in the kernel, every $g\boldsymbol{k}$ is there too. Hence $\operatorname{Ker}\varphi$ is a vector subspace of U (we know that from linear algebra) that also knows how $\mathbb{C}G$ acts on it (it's a subset of U and U knows it), and no matter what element of $\mathbb{C}G$ we apply, the result still belongs to the kernel. It cannot ever leak out of it. Hence *the kernel of any* $\mathbb{C}G$*-homomorphism is a* ***submodule*** *of* U.

[6] In words, the second condition means that any element $g \in \mathbb{C}G$ acts on the "source" module U in the same way as it acts on the "result" module V, which is what we need.

What about the image? Let's say that some vector $\boldsymbol{v} \in V$ belongs to the image. This means that it can be obtained as a result of applying the homomorphism φ on some $\boldsymbol{u} \in U$, so $\boldsymbol{v} = \varphi(\boldsymbol{u})$. However, if we apply φ on $g\boldsymbol{u}$ (for any $g \in \mathbb{C}G$), we get $\varphi(g\boldsymbol{u}) = g\varphi(\boldsymbol{u}) = g\boldsymbol{v}$, because φ is a $\mathbb{C}G$-homomorphism. So every $g\boldsymbol{v}$ is in the image as well, and *the image of any $\mathbb{C}G$-homomorphism is a submodule of V.*

6 Schur's lemma

Perhaps you're now asking "fine, $\operatorname{Ker}\varphi$ and $\operatorname{Im}\varphi$ are submodules. So what?" But we really should be interested in that, since we want to study irreducible modules (i. e. modules without submodules), and these simple facts already put stringent conditions on maps between such modules.

So suppose that we have two *irreducible* $\mathbb{C}G$-modules, U and V, and a $\mathbb{C}G$-homomorphism φ between them. What does φ look like? There's not really much freedom here, because $\operatorname{Ker}\varphi$ must be a submodule of U, and since U is irreducible, it has only the two trivial submodules: the zero and U itself. In the first case, φ must be a bijection[7], and the image is all of V. In the second case, φ maps everything to zero, and the image is the zero module. Nothing else is possible.

This gives the following: *Any $\mathbb{C}G$-homomorphism between two irreducible $\mathbb{C}G$-modules is either a bijection, or a zero map.* This result is called the *Schur's lemma* and though it is simple, it is powerful enough to serve as a base out of which everything in the representation theory may be derived.

There is also a related question: what $\mathbb{C}G$-homomorphisms can exist from an irreducible module U back to itself? Such a linear map $\varphi : U \to U$ can be represented by a square matrix M (if we pick some basis in U). This matrix has some eigenvalues given by the equation $\det(M - \lambda E) = 0$. We pick one of them and consider the map $\varphi - \lambda$. It is also a $\mathbb{C}G$-homomorphism from U to U, so Schur's lemma may be applied to it. Hence $\varphi - \lambda$ must be a bijection, or zero. But its determinant is zero (λ was picked in such a way that it would be!), so it cannot be a bijection. Hence $\varphi - \lambda = 0$ and *the only possible $\mathbb{C}G$-homomorphisms from an irreducible module to itself is a multiplication by a constant.*[8] This result is often also called the Schur's lemma and it too will be very useful for us in the following.

7 More consequences of Schur's lemma

I told you that Schur's lemma was powerful enough to serve as a foundation for the whole representation theory, so it would be good to start deriving some of its consequences. Before we do that, I'll summarize the results of the previous section in different words: *A $\mathbb{C}G$-homomorphism between two irreducible $\mathbb{C}G$-modules is just a scalar multiplication if the two modules are isomorphic* (often we would just say "the same")*, and zero map otherwise.*

Now assume that we have a general (perhaps reducible) $\mathbb{C}G$-module U and an irreducible $\mathbb{C}G$-module V, and we want to know how $\mathbb{C}G$-homomorphisms from U to V may look like. This is easy to find out, since U can be decomposed as a direct sum of irreducible modules due to Maschke's theorem, and we can write

$$U = U_1 \oplus U_2 \oplus \cdots \oplus U_n,$$

where each of the summands is an irreducible $\mathbb{C}G$-module. This means that no matter what element of $\mathbb{C}G$ is applied, vectors in each U_k will stay in it. Hence to construct a homomorphism, it suffices to decide how to map each U_k by itself. When this is known, we can map any vector $\boldsymbol{u} \in U$ by breaking it into a sum of $\boldsymbol{u}_1$ from U_1, $\boldsymbol{u}_2$ from U_2, etc., and then (because the homomorphism must be linear) do

$$\varphi(\boldsymbol{u}) = \varphi(\boldsymbol{u}_1 + \boldsymbol{u}_2 + \cdots + \boldsymbol{u}_n) = \varphi(\boldsymbol{u}_1) + \varphi(\boldsymbol{u}_2) + \cdots + \varphi(\boldsymbol{u}_n).$$

[7] Each $\boldsymbol{u} \in U$ is mapped to something in V, and each to something different. If two $\boldsymbol{u}_1 \neq \boldsymbol{u}_2$ map to the same $\in V$, their difference must map to zero due to linearity, and so it must be in the kernel. But the kernel has only the zero vector in it.

[8] Here it is crucial that everything we do is over complex numbers. Then $\det(M - \lambda E)$ is a polynomial in λ with complex coefficients, so it always must have a complex root λ. In real numbers, for instance, this would not work (for instance, the real matrix $\left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$ does not have a real eigenvalue)! However, as physicists, we love complex numbers and we use them everywhere, so we don't need to worry about this result being true only when complex numbers are allowed. (In fact there are other "algebraically closed fields" that have this property. But we should leave them to the mathematicians.)

But mapping each of the irreducible parts is easy. According to what we said above, we see that each part not isomorphic to V must map to zero. For each part that is isomorphic to V, all possible homomorphisms are given just by multiplication by a (complex) constant. So if we want to describe all possible $\mathbb{C}G$-homomorphisms from U to V, we must give one constant for each summand in U that is isomorphic to V.

To summarize: *if an irreducible $\mathbb{C}G$-module V is not present in U, then the only $\mathbb{C}G$-homomorphism possible from U to V is the zero map. If V is present n times in U, all possible $\mathbb{C}G$-homomorphisms from U to V are given by n complex constants.* As we can see, this can be a good basic tool for finding out how many times is an irreducible representation present in some module — we just have to find all possible homomorphisms between them and count them.

8 Regular module

We have already proved some nice general results about modules and homomorphisms. However, the task that we originally set out to solve is to find, for instance, what kinds of irreducible modules are at all *possible* for a given group. And we haven't made any progress with that. To make contact between groups and modules, it would be very useful to come up with a way of embodying properties of any given group in a module.

Fortunately it is quite easy to do that. We just have to remember our group algebra $\mathbb{C}G$ and the fact that we can form products of elements of $\mathbb{C}G$; of course, the group algebra is also a vector space. Thus the group algebra $\mathbb{C}G$ is a vector space that can be acted upon by elements of $\mathbb{C}G$, and so it is itself a $\mathbb{C}G$-module. This module is called *the regular module* of a given group G.

Just to put it in another perspective: in general $\mathbb{C}G$-modules, we know how to apply elements of group algebra $\mathbb{C}G$ to vectors of some vector space. In the regular module, we apply elements of $\mathbb{C}G$ to $\mathbb{C}G$ itself. Similarly, a general $\mathbb{C}G$-homomorphism $\varphi : U \to V$ must respect the group operation, so for any $g \in \mathbb{C}G$ and $\boldsymbol{u} \in U$, we must have $\varphi(g\boldsymbol{u}) = g\varphi(\boldsymbol{u})$. But if the source module is the regular module, the vectors $\boldsymbol{u} \in U$ are in fact elements of the group algebra too!

This gives us a powerful conclusion: if φ is a $\mathbb{C}G$-homomorphism from the regular module to some module V, we actually must have $\varphi(gh) = g\varphi(h)$ for any $g, h \in \mathbb{C}G$. We can get an even more striking result if we put $h = 1$:[9] then we have $\varphi(g1) = \varphi(g) = g\varphi(1)$ and all possible homomorphisms are just given by the choice of $\varphi(1)$!

So let's consider any irreducible $\mathbb{C}G$-module V. What are all possible $\mathbb{C}G$-homomorphisms from $\mathbb{C}G$ to V? The only choice we have is the vector of V that we pick to be the result of $\varphi(1)$. This means that all such possible homomorphisms are given by $\dim V$ complex constants (the components of the vector $\varphi(1) \in V$ in some arbitrary basis).

Hence we see that not only *each irreducible $\mathbb{C}G$-module that can exist is present in the regular module*, but *the number of times it appears in the regular module is equal to its dimension*! This means that the number of possible irreducible $\mathbb{C}G$-modules for a given group is severely limited — they all must "fit" into the regular module. In fact, since the regular module can be written as a direct sum of these irreducible modules (courtesy of Maschke's theorem again!), their dimensions must exactly add up to the dimension of the regular module, which is $|G|$.[10]

Let's denote U_1, U_2,..., U_n all the different (i. e. non-isomorphic) irreducible modules present in $\mathbb{C}G$ (and we know that these are all that are possible!). Then U_1 is present $\dim U_1$ times, U_2 is present $\dim U_2$ times etc. In the end, we have

$$|G| = \sum_k (\dim U_k)^2.$$

That's already a bunch of pretty powerful results. And we'll find more!

[9]By "1" I mean $1e$, or one times the identity element of the group. This is perfectly in line with what lazy physicists do — they conflate numbers with multiples of identity operators all the time.

[10]Remember, it's just $\mathbb{C}G$ — all formal linear combinations of group elements. So the $|G|$ group elements form a natural basis of $\mathbb{C}G$.

9 Decomposition of the regular module

Let's now decompose $\mathbb{C}G$ in a slightly different manner. We know that each irreducible module U_k is present $\dim U_k$ times in $\mathbb{C}G$, so we can lump them together and denote the direct sum of all these copies by W_k. Then we have

$$\mathbb{C}G = W_1 \oplus W_2 \oplus \cdots \oplus W_n,$$

where each W_k is a direct sum of all the copies of the irreducible module U_k. Of course, this means that if we pick two different W's they do not share any irreducible submodules with each other (since each is built out of copies of a different irreducible module). This will give us more powerful results.

Let's say that we pick some $w_1 \in W_1$ and multiply by it all vectors of $\mathbb{C}G$ *from the right*. Each $g \in \mathbb{C}G$ gets sent to gw_1. But here's the catch: W_1 is a *submodule* of $\mathbb{C}G$, which means that no matter what group element is used to act on it, the result stays in the submodule. So any gw_1 will in fact lie in W_1.

Now let's ask what happens to a vector $w_2 \in W_2$ when we form this product. It goes to $w_2 w_1 \in W_1$. We also notice that this mapping is trivially a $\mathbb{C}G$-homomorphism[11] from W_2 to W_1, and since W_1 and W_2 do not share any equivalent irreducible submodules, the only possible $\mathbb{C}G$-homomorphism between them is zero. This means that *if we multiply any vectors from different W's, the product is always zero*!

And we can get more. Since $\mathbb{C}G$ can be written as a *direct* sum of the W's, it means that *any* vector in $\mathbb{C}G$ can be resolved into a sum of $w_1 + w_2 + \cdots + w_n$, where each w_k comes from its own W_k, and this can be done in one way only. In particular, this is also true for the unit element 1 of the algebra $\mathbb{C}G$, so we have

$$1 = e_1 + e_2 + \cdots + e_n, \quad \text{where each } e_k \in W_k,$$

and this defines the e_k's precisely, without any freedom. Since each e_k comes from a different W_k it is still true that a product of two different e_k's is zero. But they have another nice property which can be obtained by a simple trick. For any vector w_k from some W_k, we can write the following:

$$w_k = 1w_k = (e_1 + e_2 + \cdots + e_n)w_k = e_k w_k,$$

since all other e's belong to a different W, and the product of any two vectors from different W's is zero. What's more, it also works from the other side:

$$w_k = w_k 1 = w_k(e_1 + e_2 + \cdots + e_n) = w_k e_k.$$

So if we multiply any vector g by e_k from either side, the part of g that belongs to W_k is kept untouched, and everything else is killed. In particular, e_k belongs to W_k, so if we multiply it by e_k, it will be untouched. Hence $e_k^2 = e_k$.

It is clear that these e_k are quite important, and they deserve to get a special name. We call them the *idempotents*.[12] Let us summarize their nice properties:

1. Each idempotent is a projection operator onto its corresponding W (i. e. it keeps the W-part of any vector alone and kills the rest).
2. Hence $e_k^2 = e_k$, and any power of the idempotent is equal to the idempotent itself.
3. And if we multiply two different idempotents, each tries to project the other into its own space and the result must be zero.

But remember: **all of this works only if the W's do not share irreducible submodules!** Otherwise we cannot ensure that $w_k w_\ell = 0$, and the whole argument crumbles.

10 Characters

We already know a bit about characters (from the text about modes of the square), but it will certainly be good to repeat the important facts once more. We'll also need to translate the notion of the character to our new language of modules etc.

[11] Because if $\varphi(w) = ww_1$, then we must have $\varphi(gw) = g\varphi(w)$, which obviously works.

[12] "idem-", not "im-"! In Latin, "idem" means "the same". The idea is that the powers of e_k are *the same* as e_k itself.

Fortunately, this is not too difficult. We already introduced characters as traces of representation matrices, and in our treatment, the matrices have been replaced by the modules. In the matrix language, the emphasis is on the tables of numbers. For us, it is on the underlying vector spaces that know what to do if acted upon by elements of a group — the modules.

The module is (by definition) a vector space on which group elements act linearly, so multiplication by any group element corresponds to a certain linear mapping from the module back to itself. This means that in any basis that we pick, the mapping is described by a square matrix, and the trace of the matrix is independent of the choice of the basis. Hence these traces are the property of the module itself (of course, for each group element the trace may be different).

Since we already talked a bit about characters in the document about the modes of the square, we just repeat the important properties without proof here:

- $\chi(1)$ is always equal to the dimension of the module.
- All elements in the same conjugacy class have the same character.
- For a k-dimensional module, the matrices are $k \times k$. Say that g is an element of order n, so that $g^n = 1$. The same holds for the matrices too, so their eigenvalues must be n-th roots of unity. Each character is then a sum of k such numbers.
- The element g^{-1} has a matrix whose eigenvalues are all reciprocal to the original ones. But, since the eigenvalues are roots of unity, their reciprocals are the same as their complex conjugates. Hence $\chi(g^{-1}) = \chi(g)^\star$.

11 Orthogonality

In this section, we get another powerful result: it turns out that we will be able to calculate the idempotents just by looking at the character table, and from the idempotents, we will be able to prove nice properties of the characters themselves.

However, if we want to make contact with the characters, we first have to find a way of obtaining them using our present formalism. Fortunately this is not too hard. If we want to get the character of the module W_k, we just have to consider what happens if any $g \in G$ (just in the group, not in the whole algebra!) acts on it. This will take any $w_k \in W_k$ to a product gw_k which is still in W_k, so this is a linear map $W_k \to W_k$. Its trace is exactly the character of the element g in W_k. But this character is not exactly the one that we look up in the table, because the table shows characters of *irreducible* modules. And each W_k is made of multiple copies of the same irreducible module. This means that in a nice basis, the matrices of the map $w_k \to gw_k$ are all block-diagonal; the blocks are all the same. This is because W_k is a sum of some number of equivalent irreducible modules. Let's call that number n_k. Then we have n_k blocks on the diagonal, all of them are the same and each of them has the trace $\chi_k(g)$ that we can look up in the table. Hence the character of g in W_k is $n_k\chi_k(g)$.

We can also obtain this number by acting on the whole $\mathbb{C}G$. We just have to first project down to W_k. So if we consider the map that takes any $r \in \mathbb{C}G$ and maps it to ge_kr, we find that its trace is also $n_k\chi_k(g)$ (because e_k discards everything that does not belong to W_k).

Now let's consider the same map from a different point of view. e_k is an element of $\mathbb{C}G$, so it must be a certain linear combination of the group elements, and we can write $e_k = \sum_{h \in G} c_h h$. Can we find the trace of ge_kr independently? Yes: we can replace e_k by the sum and find that the result is $\sum_{h \in G} c_h ghr$. Now we can calculate the trace easily, because we already know that the trace of any map $r \to gr$ is $|G|$ if $g = 1$ and zero otherwise. This means that the trace is just $|G|c_{g^{-1}}$ (because $h = g^{-1}$ is the only element that makes $gh = 1$), and we have

$$|G|c_{g^{-1}} = n_k\chi_k(g) \qquad \Longrightarrow \qquad c_g = \frac{n_k\chi_k(g^{-1})}{|G|},$$

so the idempotent may be written as

$$e_k = \frac{n_k}{|G|}\sum_{g \in G}\chi_k(g^{-1})g = \frac{n_k}{|G|}\sum_{g \in G}\chi_k^\star(g)g.$$

But we know that $e_k^2 = e_k$. Plugging the explicit expression above into it, we get

$$\frac{n_k^2}{|G|^2} \sum_{g \in G} \sum_{h \in G} \chi_k(g^{-1}) \chi_k(h^{-1}) gh = \frac{n_k}{|G|} \sum_{g \in G} \chi_k(g^{-1}) g.$$

Let's find out the coefficient at 1 on both sides. In the sum on the left, we pick only the terms with $gh = 1$, or $h^{-1} = g$. This gives

$$\frac{n_k^2}{|G|^2} \sum_{g \in G} \chi_k(g^{-1}) \chi_k(g) = \frac{n_k}{|G|} \chi_k(1) \qquad \Longrightarrow \qquad \frac{1}{|G|} \sum_{g \in G} \chi_k(g) \chi_k^\star(g) = 1.$$

On the other hand, $e_k e_\ell = 0$. Doing the same trick, we find

$$\frac{n_k n_\ell}{|G|^2} \sum_{g \in G} \sum_{h \in G} \chi_k(g^{-1}) \chi_\ell(h^{-1}) gh = 0,$$

or after canceling $n_k n_\ell / |G|$ and picking only the coefficient at 1,

$$\frac{1}{|G|} \sum_{g \in G} \chi_k^\star(g) \chi_\ell(g) = 0.$$

So if we introduce a scalar product of two functions $\varphi(g)$, $\psi(g)$ (that take a group element g and yield a complex number) like this:

$$\langle \varphi | \psi \rangle = \frac{1}{|G|} \sum_{g \in G} \varphi^\star(g) \psi(g), \tag{1}$$

we find that the irreducible characters (the rows in the character table) are orthonormal:

$$\langle \chi_k | \chi_\ell \rangle = \delta_{k\ell}. \tag{2}$$

12 Class functions

We already noted that the characters are always the same for all elements of the same conjugacy class, and so they are really just functions of conjugacy classes. Any such function is called a *class function*. Those functions can be fully specified just by giving one complex number for each conjugacy class (the result of the function of that class). We can also add them together and multiply them by scalars. This means that class functions form a k-dimensional vector space, where k is the number of conjugacy classes.

The characters of irreducible modules are, of course, class functions, and according to (2), they are orthonormal with respect to the scalar product in (1), which means that they must be linearly independent. From that alone, we immediately see that there can be *at most* as many irreducible characters as there are classes.

That's quite an interesting result. Can we make it more precise? It turns out that we can if we look at the group algebra $\mathbb{C}G$ once again. Let's say that the group has conjugacy classes K_1, K_2, $\ldots K_k$. For each of them, we form the sum of all the elements in it:

$$s_\ell = \sum_{h \in K_\ell} h.$$

These s_1, s_2, etc. are, of course, elements of the group algebra. However, they are quite special: since each class contains all elements that can be obtained one from another by conjugation, it follows that if we conjugate all elements in a class by a single element g, we obtain all the elements of the class again (just maybe in a different order). Hence, if we conjugate any s_ℓ, the terms in the sum can shuffle around, but we get the same sum again:

$$g s_\ell g^{-1} = s_\ell \quad \text{for any } g \in G \text{ and any } s_\ell.$$

This means that each such s_ℓ commutes with every group element, and so it commutes with everything in the whole group algebra itself. This is very useful, because it means that left multiplication by any s_ℓ is a homomorphism. Technically said, if we define $\varphi(g)$ to be $s_\ell g$, we see that

$$\varphi(gh) = s_\ell gh = g s_\ell h = g\varphi(h),$$

and φ truly is a homomorphism.

Once more, we utilize Schur's lemma. Consider an irreducible module U. Since it's a module, multiplication by any s_k always ends up in U, and so s_k is a homomorphism from U to U. Hence it acts on U just as a multiplication by a constant. The same is true for the submodules W_k that contain all the copies of a single irreducible module (multiplication by s_k acts as multiplication by the same constant on all the copies). Hence we can write

$$s_k = s_k 1 = s_k(e_1 + e_2 + \cdots + e_n) = \lambda_1 e_1 + \lambda_2 e_2 + \cdots + \lambda_n e_n,$$

where the e's are the idempotents (each of them comes from its own W) and the λ's are some complex constants.

Hence each s_k can be written as a linear combination of the idempotents, and (since each s_k is a sum of different group elements) all of them are linearly independent. Hence there are at most as many classes (s_k's) as there are irreducible characters (idempotents). Putting it together with the previous result, we see that *the number of irreducible modules possible for the given group G is always equal to the number of conjugacy classes of G.*

Since the irreducible characters are linearly independent class functions, and there are as many of them as there are classes, we see that the irreducible characters form a basis of the vector space of all class functions. This means that *any class function can be written as a linear combination of the irreducible characters.*

One nice class function that can be written in this manner is the *indicator function* of a given class: the function that gives 1 for members of the class and 0 for everything else:

$$\Delta_\ell(g) = \begin{cases} 1 & \text{if } g \text{ belongs to the class } K_\ell, \\ 0 & \text{otherwise.} \end{cases}$$

We write it as a linear combination of irreducible characters:

$$\Delta_\ell(g) = \sum_{p=1}^{n} c_{p,\ell} \chi_p(g),$$

where $c_{p,\ell}$ are some unknown coefficients. These can be found easily if we use the orthogonality conditions (2): if we take the scalar product of both sides with one irreducible character χ_r, we simply find that $c_{r,\ell} = \langle \chi_r, \Delta_\ell \rangle = \frac{1}{|G|} \sum_{g\in G} \chi_r(g)^\star \Delta_\ell(g)$. The indicator function $\Delta_\ell(g)$ is zero for all elements that do not belong to the class K_ℓ, and for those that do, it gives one. Hence we sum $\chi_r(g)^\star$ for each element of the class, and we get

$$\Delta_\ell(g) = \sum_p \frac{|K_\ell|}{|G|} \chi_p(K_\ell)^\star \chi_p(g),$$

where $|K_\ell|$ is just the size of the class, and $\chi_p(K_\ell)$ is the character of any of its elements. This can also be rewritten as

$$\sum_p \chi_p(g)^\star \chi_p(h) = \begin{cases} |G|/|K_\ell| & \text{if both } g \text{ and } h \text{ belong to the same class } K_\ell, \\ 0 & \text{if they belong to different classes.} \end{cases}$$

If you think about it a little bit, you will find that this means that *the columns of the character table are also orthogonal.*