Modules I: Basic definitions and constructions

October 23, 2014

1 Definition and examples

Let R be a ring. A *left R-module* is an abelian group M equipped with a map $\theta : R \times M \longrightarrow M$ satisfying the three conditions listed below; we usually omit θ from the notation and write simply rx or $r \cdot x$ for $\theta(r, x)$:

1. distributive law: $(r_1 + r_2)x = r_1x + r_2x$ and $r(x_1 + x_2) = rx_1 + rx_2$.
2. associative law: $r(sx) = (rs)x$.
3. identity: $1 \cdot x = x$ for all $x \in M$.

Right R-modules are defined similarly; we put the "r" on the right so that the associative law is $(xr)s = x(rs)$. As with the group actions, the distinction is just that in a left module rs acts by s first, then r; in a right module r acts first and then s.

1.1 On left versus right modules, and the opposite ring

In constrast to group actions, however, one cannot always convert left module structures to right structures and vice versa. To clarify the situation, it is convenient to introduce the *opposite ring* R^{op}, which has the same underlying abelian group R but with multiplication $a \cdot b = ba$. Then a left module over R^{op} is the same thing as a right module over R, and vice versa. Consequently, if R^{op} happens to be isomorphic to R—preferably in some canonical way—then we can convert back and forth between left and right actions in much the same way as we did for group actions. Explicitly, suppose we are given an isomorphism $\alpha : R \xrightarrow{\cong} R^{op}$, and M is a right R-module. Then $r \star x = x\alpha(r)$ $(r \in R, x \in M)$ makes M into a left R-module. Note also that such an isomorphism α is the same thing as an *anti-automorphism* of R, i.e. an isomorphism of abelian groups $\alpha : R \xrightarrow{\cong} R$ such that $\alpha(ab) = \alpha(b)\alpha(a)$ for all $a, b \in R$.

Many familiar rings admit such anti-automorphisms:

- If K is a commutative ring and $R = M_nK$, then $A \mapsto A^T$ (the transpose) is an anti-automorphism.
- If K is a commutative ring, G is a group and $R = KG$, then $g \mapsto g^{-1}$ (extended K-linearly) is an anti-automorphism.

- If $x = a + bi + cj + dk$ is a quaternion, its *quaternionic conjugate* is defined by $\overline{x} = a - bi - cj - dk$. Quaternionic conjugation is an anti-automorphism of $\mathbb{H}$, as can be checked by direct computation.

- If R is a commutative ring, then the identity is both an automorphism and an anti-automorphism.

Thus in all of these examples, one has the option of converting between left and right modules. As in the case of G-sets, however, it is still essential to pay close attention to which side the ring is acting on (well, except in the commutative ring example, where left and right modules are the same thing). For an example of a ring which is not isomorphic to its opposite, i.e. which does not admit any anti-automorphism, see the exercises.

1.2 An alternate view of modules

This section is exactly analogous to the one entitled "An alternate view of group actions" in the G-set notes. Recall that in that section we showed how to interpret a left group action $G \times S \longrightarrow S$ can be interpreted as a group homomorphism $G \longrightarrow Perm\, S$. Similarly, a left R-module structure $\theta : R \times M \longrightarrow M$ is the same thing as a ring homomorphism $\phi : R \longrightarrow End_{\mathbb{Z}} M$: Given θ, for $r \in R$ we let $\phi(r)(x) = rx$. Given ϕ, we define rx by the same equation. One can easily check that this works. If M is a right module, one obtains similarly a homomorphism $R \longrightarrow (End_{\mathbb{Z}} M)^{op}$ (check this!—you'll find that the order of composition gets reversed, hence the "op").

If R is an F-algebra, then M is an F-vector space. In this case we get a homomorphism of F-algebras $\phi : R \longrightarrow End_F M$. Thus for each $r \in R$ the endomorphism $\phi(r)$ constructed in the previous paragraph is not only an abelian group homomorphism but an F-linear map. For example, in the case of a group representation, i.e. an FG-module V, we get an F-algebra homomorphism $FG \longrightarrow End_F V$. This is the same homomorphism defined in the "Algebras over a field" notes.

1.3 The category of R-modules, and first examples

If M, N are left R-modules, an *R-module homomorphism* $\phi : M \longrightarrow N$ is a homomorphism of abelian groups such that $\phi(rx) = r\phi(x)$ for all $r \in R$, $x \in M$. Since identity maps are R-module homomorphisms and the composition of two R-module homomorphisms is an R-module homomorphism, left R-modules form a category that we denote **R-mod**. Similarly, there is a category **Mod-R** of right R-modules. Some examples:

- If F is a field, an F-module is the same thing as a vector space over F. Module homomorphisms are the same thing as F-linear maps; thus **F-mod**=**F-Vect**.

- A $\mathbb{Z}$-module is the same thing as an abelian group. Here $n \cdot x = x + x + ... + x$ (n terms). Module homomorphisms are the same thing as group homomorphisms; thus $\mathbb{Z}$**-Mod**=**Ab**.

- If I is a left ideal in R, then I is a left R-module. Similarly for right ideals. In particular, R itself is both a left and a right R-module under its ring multiplication.

- If F is a field, the usual multiplication of column vectors by matrices $(A, X) \mapsto AX$ ($A \in M_nF$, $X \in F^n$) makes F^n a left module over M_nF. Similarly right multiplication $(X, A) \mapsto XA$, where X is now a row vector, defines a right module structure.

- For any ring R, the *zero module* is the zero group $\{0\}$ with its unique R-module structure.

The next two examples deserve sections of their own.

1.4 Modules over a polynomial ring

Consider the polynomial ring $F[x]$, and let V be an $F[x]$-module. Then V is an F-vector space and the module structure is equivalent to giving an F-algebra homomorphism $\rho : F[x] \longrightarrow End_F V$. By the universal property of polynomial rings, this in turn is equivalent to just specifying one element $A \in End_F V$, namely $A = \rho(x)$. So if we start with A, we get an $F[x]$-module we'll denote V_A.

Now assume that $dim_F V = n < \infty$. Then linear algebra, one of the most widely applied subjects in mathematics (applied both within pure mathematics and in the sciences, business, etc.) can be viewed as the study of such transformations A (along with linear transformations between two different vector spaces). Now we see that large parts of linear algebra become a chapter in module theory. For example, one of the main jobs of the linear algebraist is to find a basis for V in which the matrix of A takes a particularly simple form: triangular, diagonal, etc. Let's see how this translates into the $F[x]$-module world.

In basic linear algebra the question is usually phrased in terms of "similar" matrices. Here we assume a fixed basis for V has been chosen, so we might as well take $V = F^n$ with the standard basis. Then A, B are said to be *similar* if there is an invertible matrix P with $B = PAP^{-1}$. This is equivalent to saying that we have a new basis, namely $\{v_i := Pe_i (1 \leq i \leq n)\}$, such that the matrix of linear transformation A in the new basis is B. For example, if we were trying to show A is diagonalizable (i.e. there is a basis of eigenvectors) we would need a P as above with B a diagonal matrix. Now, rather than just finding a good basis for one linear transformation, we might want to classify all matrices up to similarity; or in our G-set terminology, determine the orbits of the conjugation action of GL_nF on M_nF.

Proposition 1.1 *A and B are similar if and only if V_A and V_B are isomorphic as $F[x]$-modules.*

Proof: We denote the V_A module structure using a dot $\cdot$, and the V_B structure using a star $\star$. Suppose $B = PAP^{-1}$. Here P is an F-linear isomorphism $V \longrightarrow V$. On the other hand

$$P(x \cdot v) = PAv = BPv = x \star P(v).$$

Hence P is an $F[x]$-module homomorphism and therefore an $F[x]$-module isomorphism. For the converse just run this argument in reverse: If we are given an isomorphism of modules $P : V_A \xrightarrow{\cong} V_B$, then P is in particular an F-linear isomorphism. Using the same equations as above we find that $PA = BP$ as desired.

Thus our linear algebra problem has been converted to a problem about modules over the principal ideal domain $F[x]$. As we will see in subsequent installments of the notes, this problem has a complete, elegant solution.

1.5 Representations of groups and modules over the group algebra

Representations of groups were invented and studied before modules, and before the abstract concept of a ring was fully developed, so suppose for a moment that the concepts of "module" and "group algebra" have not yet been invented. Let V be a representation of G over F of dimension n, where for illustrative purposes we will take G finite and $F = \mathbb{R}$. Choosing a basis for V for concreteness, we think of our representation as a homomorphism $\rho : G \longrightarrow GL_nF$. Since $GL_nF \subset M_nF$, we are free to take linear combinations of the elements $\rho(g)$ inside M_nF. And this is something one definitely wants to do. For example, one might want to know the invariants $V^G = \{v \in V : gv = v \, \forall g \in G\}$. Since $\mathbb{R}$ has characteristic zero, we can define an averaging operator

$$f := f_\rho = \frac{1}{|G|} \sum_{g \in G} \rho(g) \in M_nF.$$

Then by a familiar calculation one checks that f is idempotent and in fact is a projection onto V^G, so in particular $V^G = Im\, f$. The unsatisfying thing is that f depends on the particular ρ, whereas the averaging operator longs for a home of its own, depending only on the group and not on the representation.

This leads us to define the group algebra FG—$\mathbb{R}G$ in this case—which provides a happy home indeed for our universal averaging operator e_0, defined in the exercises to "Algebras over a field". Upon further reflection (we are compressing many years of mathematical history into a moment) we realize that not only does G act on V but in some sense the entire group algebra FG "acts" on V. Pondering what it might mean for an algebra (or even just a ring) such as FG to "act" on V, we arrive at the definition of a module structure $FG \times V \longrightarrow V$.

Returning to the present where we have in fact defined group algebras and modules, the upshot of the discussion is that a representation of G over F is the same thing as an FG-module. Explicitly, if we have an FG-module V with scalar multiplication $FG \times V \longrightarrow V$, by simply restricting to $G \times V \subset FG \times V$ we get a linear action of G on V, i.e. a representation. Conversely if we start with a linear action $G \times V \longrightarrow V$, we extend it by the distributive law to $FG \times V \longrightarrow V$ and we have our FG-module. Alternatively, one can think in terms of the correspondence $Hom_{grp}(G, GL(V)) \leftrightarrow Hom_{F-alg}(FG, End_F V)$ described earlier. Then if $\rho : G \longrightarrow GL(V)$ and $\phi : FG \longrightarrow End_F V$ is the corresponding algebra homomorphism, $\phi(e_0) = f_\rho$. Thus we indeed have one averaging operator e_0 that efficiently determines all of the old-fashioned f_ρ's at once.

How can we actually construct representations of a group? One easy way is via *permutation representations*. If X is any set, we let FX denote the vector space with basis X. Thus an element of FX is a formal linear combination $\sum_{i=1}^n a_i x_i$ with $a_i \in F$, $x_i \in X$. Note this agrees with our notation FG. If X is a left G-set, then by extending the action map

$G \times X \longrightarrow X$ linearly over sums, we get a linear G-action on FX: $g \cdot \sum a_i x_i = \sum a_i g x_i$, and hence an FG-module. Such representations are called *permutation representations*.

Some examples:

- Take $X = G$ with G acting by left translation. This is called the *regular representation*. In terms of modules, it is the left regular module of FG.
- More generally, take $X = G/H$ for some subgroup H, with the left translation action.
- Let $G = S_n$ and $X = [n]$, with the standard action. Then $FX = F^n$ with S_n acting by permuting the standard basis vectors; in other words, via the standard inclusion $S_n \subset GL_n F$. This is also a special case of the previous item, with $H = S_{n-1}$.

This construction defines a functor **G-Set** $\longrightarrow$ **FG-Mod**. As usual this is a lie if taken literally, since I haven't defined the map of modules induced by a map of G-sets. Also as usual, the required definition is obvious, and checking the functor axioms is trivial.

2 Some basic constructions

Submodules. Submodules are defined in the obvious way: $N \subset M$ is a submodule if it is a subgroup and it is closed under scalar multiplication: for all $r \in R$ and $x \in N$, $rx \in N$. For vector spaces this is the same thing as a vector subspace in the linear algebra sense, and for abelian groups a submodule is just a subgroup. A left ideal $I \subset R$ is by definition a submodule of R (where the latter is regarded as a left module over itself).

Example. Consider the ring of matrices $M_n F$, F a field, regarded as a left module over itself. For $1 \leq i \leq n$, let C_i denote the matrices with j-th column zero for $j \neq i$. Then C_i is a left ideal and hence an $M_n F$-submodule.

Quotient modules. Suppose $N \subset M$ is a submodule. In particular N is a subgroup of the abelian group M, so we can form the quotient group M/N. We then define a multiplication by $r \cdot (x + N) = rx + N$. This is a well-defined R-module structure, and the quotient map $\pi : M \longrightarrow M/N$ is a surjective R-module homomorphism. For example, given any left ideal $I \subset R$ we can form the quotient module R/I. For $\mathbb{Z}$-modules, the quotient module is just the quotient group.

The quotient for vector spaces usually isn't discussed in undergraduate linear algebra courses, but is easy to work with. For example, suppose V is an F-vector space, of finite dimension n for simplicity, and W a linear subspace ($=F$-submodule) of dimension k. Choose a basis $v_1, ..., v_n$ for V such that $v_1, ..., v_k$ is a basis for W. Then the images of $v_{k+1}, ..., v_n$ in V/W form a basis for V/W.

Products and direct sums. Let M_α be a collection of R-modules, $\alpha \in J$ for some index set J. Then the product set $\prod_\alpha M_\alpha$ has a natural R-module structure, with addition and scalar multiplication given component-wise; thus each projection $\pi_\beta : \prod M_\alpha \longrightarrow M_\beta$ is an R-module homomorphism. In fact this is the categorical product: If N is an R-module

and $f_\alpha : N \longrightarrow M_\alpha$ a collection of R-module homomorphisms, there is a unique R-module homomorphism $f : N \longrightarrow \prod_\alpha M_\alpha$ such that $\pi_\alpha \circ f = f_\alpha$.

The *direct sum* or *coproduct* of the M_α's is the submodule $\oplus_\alpha M_\alpha$ of the product consisting of elements having all but finitely many components equal to zero. This is in fact the categorical coproduct: If N is an R-module and $f_\alpha : M_\alpha \longrightarrow N$ a collection of R-module homomorphisms, there is a unique R-module homomorphism $f : \oplus_\alpha M_\alpha \longrightarrow N$ such that $f \circ i_\beta = f_\beta$ for all $\beta \in J$, where $i_\beta : M_\beta \longrightarrow \oplus_\alpha M_\alpha$ is the evident inclusion. Indeed, we just take $f(\{x_\alpha\}) = \sum f_\alpha(x_\alpha)$, which makes sense since only finitely many of the x_α's are nonzero.

If the index set is finite then direct sum and product coincide. But in the infinite case they exhibit markedly different behaviour. For example, let $\oplus_{i=1}^\infty M$, $\prod_{i=1}^\infty M$ denote the direct sum and direct product of a countably infinite number of copies of the same module M. Then if M is nonzero, the product will always have strictly greater cardinality than the direct sum; e.g. if $R = M = \mathbb{Z}$ then the direct sum is countable but the product is not. Or consider the case of abelian groups. Any direct sum of abelian torsion groups (i.e. groups in which every element has finite order) is again a torsion group. But this is false for infinite products; consider $\prod_{i=1}^\infty \mathbb{Z}/p^i$. The element $(1, 1, 1, ...)$ has infinite order.

Generators. Let X be a subset of a module M. The *submodule generated by* X is the smallest submodule of M containing X, denoted $\langle X \rangle$ or RX. (*Caution:* The latter notation is also used for free modules; see below.) Equivalently $\langle S \rangle$ is the set of all elements of the form $\sum r_i x_i$, $x_i \in S$ and $r_i \in R$.

For $R = \mathbb{Z}$ this is just the usual notion of generators for abelian groups. For vector spaces it is the same thing as the "span" of the set in linear algebra. For left ideals it is the same thing as "left ideal generated by the set X", a concept that predates the invention of modules.

Cyclic modules and annihilator ideals. If $\langle S \rangle = M$ we say that S generates M. An R-module is *cyclic* if it is generated by a single element. If $x \in M$ is a generator, then the map $\eta : R \longrightarrow M$ given by $\eta(r) = rx$ is an R-module map whose kernel is a left ideal denoted $ann\, x$ and called the *annihilator ideal* of x. Thus $ann\, x = \{r \in R : rx = 0\}$, and η factors through an isomorphism $R/(ann\, x) \xrightarrow{\cong} M$. If M is any R-module, and $x \in M$, by applying this construction to the submodule generated by x we get $R/(ann\, x) \cong Rx \subset M$. For example, if $R = \mathbb{Z}$ and $x \in M$ is an element of order n, then $ann\, x = (n)$. (In fact, this is really the definition of "order n element".) If $R = M_n F$, $M = F^n$ and $x = e_1$, then $ann\, x$ is the left ideal consisting of all matrices whose first column is zero.

Exact sequences and extensions. By a "sequence" of R-modules we mean any sequence of R-module homomorphisms

$$...\longrightarrow M_{k-1} \xrightarrow{f_{k-1}} M_k \xrightarrow{f_k} M_{k+1} \longrightarrow ...,$$

which is allowed to be finite as well as infinite in one or both directions. Such a sequence is *exact* if for all k we have $Im\, f_{k-1} = Ker\, f_k$. Note the following special cases:

- $0 \longrightarrow M \longrightarrow N$ is exact if and only if $M \longrightarrow N$ is injective.

- $M \longrightarrow N \longrightarrow 0$ is exact if and only if $M \longrightarrow N$ is surjective.
- $0 \longrightarrow M \longrightarrow N \longrightarrow 0$ is exact if and only if $M \longrightarrow N$ is an isomorphism.
- $0 \longrightarrow L \xrightarrow{i} M \xrightarrow{\pi} N \longrightarrow 0$ if and only if i is injective, π is surjective, and $Im\, i = Ker\, \pi$.

The last case is called a *short exact sequence* or an *extension* of modules. It is exactly analogous to a group extension; indeed $L \longrightarrow M \longrightarrow N$ is in particular an extension of abelian groups (and indeed if $R = \mathbb{Z}$, it is the same thing as such an extension). As with group extensions, one can, if desired, regard i as the inclusion of a submodule; then N is canonically isomorphic to M/L. The one notational difference here is that with modules, we prefer to stick those zero modules at the end and call it a short exact sequence.

A short exact sequence as above is *split* if there exists a *splitting* or *section* $s : N \longrightarrow M$, i.e. a module homomorphism such that $\pi \circ s = Id_N$. In contrast to the situation for groups, a split extension is automatically a direct product (=direct sum).

Proposition 2.1 *Suppose* $0 \longrightarrow L \xrightarrow{i} M \xrightarrow{\pi} N \longrightarrow 0$ *is a short exact sequence of* R*-modules. If the sequence splits, then* $M \cong L \oplus N$ *(in a particular way to be specified in the proof).*

Proof: Let $s : N \longrightarrow M$ be a splitting. By the universal property of direct sums, there is a unique homomorphism $\phi : L \oplus N \longrightarrow M$ such that $\phi|_L = i$ and $\phi|_N = s$; i.e. $\phi(x, y) = i(x) + s(y)$. If $\phi(x, y) = 0$, then applying π to this equation we find that $y = 0$. Since i is injective, we then have $x = 0$. So ϕ is injective. Now suppose $z \in M$, and let $w = z - s\pi(z)$. Then $\pi(w) = 0$, so by exactness $w = i(x)$ for a unique $x \in L$. Hence $z = \phi(x, \pi(z))$ and ϕ is surjective.

Here's another property of splittings for modules that has no analogue for non-abelian group extensions.

Proposition 2.2 *Let* $0 \longrightarrow L \xrightarrow{i} M \xrightarrow{\pi} N \longrightarrow 0$ *be a short exact sequence of* R*-modules. Then the sequence splits if and only if there is a homomorphism* $r : M \longrightarrow L$ *such that* $ri = Id_L$.

Proof: Suppose s is a splitting, and let $q = Id_M - s\pi : M \longrightarrow M$. Then $\pi q = \pi - \pi s \pi = \pi - \pi = 0$, so $q = ir$ for a unique $r : M \longrightarrow L$ by exactness at M. To show that $ri = Id_L$, since i is injective it suffices to show $iri = i(id_L)$, or $iri = i$. But $iri = qi = (Id_M - s\pi)i = Id_M i = i$, as desired.

Conversely suppose given $r : M \longrightarrow L$ with $ri = Id_L$. Let $t = Id_M - ir$. Then $ti = i - iri = i - i = 0$. Hence $t = s\pi$ for a unique homomorphism $s : N \longrightarrow M$. To check that $\pi s = Id_N$, since π is surjective it suffices to show $\pi s \pi = \pi$. But $\pi s \pi = \pi t = \pi(Id_M - ir) = \pi Id_M = \pi$.

Examples. 1. Take $R = \mathbb{Z}$. There are many non-split short exact sequences of abelian groups, as we've already seen. $0 \longrightarrow C_2 \longrightarrow C_4 \longrightarrow C_2 \longrightarrow 0$ is the smallest, but you can cook up a million others.

2. If $R = F$ is a field, every short exact sequence splits. This can be proved directly, but the fastest way is to apply Exercise 3.

We will see many more examples soon. In particular we will determine for exactly which fields F and finite groups G it is the case that *every* short exact sequence of modules over FG splits. It is always a good thing when this happens, because then the isomorphism type of the middle term of the short exact sequence is determined by the other two modules; it is just their direct sum.

3 Free modules

Let X be a set. Then the *free R-module on X* is the direct sum over the index set X of copies of R. We denote it $RX := \oplus_{x \in X} R_x$, where R_x is the copy of R indexed by x. For example, if $X = \{1, 2, ..., n\}$ then $RX = \oplus_{i=1}^{n} R_i$. We often identify $x \in X$ with $1 \in R_x$, so there is an inclusion $i : X \longrightarrow RX$. There is the following universal property:

Proposition 3.1 *Let M be any R-module, and $f : X \longrightarrow M$ any map of sets. Then there is a unique R-module homomorphism $\phi : RX \longrightarrow M$ such that $\phi \circ i = f$.*

Proof: As seen earlier, the *left regular module* R (that is, R as a left module over itself) has the universal property: If $y \in M$, there is a unique R-module map $\psi : R \longrightarrow M$ such that $\psi(1) = y$, namely $\psi(r) = ry$. Combining this with the universal property of direct sums completes the proof.

As a consequence we see that $X \mapsto RX$ defines a functor **Set** $\longrightarrow$ **R-mod**: Given a map of sets $g : X \longrightarrow Y$, let $Rg : RX \longrightarrow RY$ be the unique R-module homomorphism such that $x \mapsto g(x)$. We then have our usual interpretations:

Plain English version: If you want to define an R-module homomorphism $RX \longrightarrow M$, it's enough (indeed equivalent) to "say where X goes", i.e. define a set map $X \longrightarrow M$.

Adjoint functor version: The free module functor **Set** $\longrightarrow$ **R-mod** given by $X \mapsto RX$ is left adjoint to the forgetful functor **R-mod** $\longrightarrow$ **Set**. In other words, there is a natural bijection

$$Hom_{R-mod}(RX, M) \cong Hom_{set}(X, M).$$

We say that M is a *free R-module* if there exists a set X such that M is isomorphic to RX. The image of X in M under such an isomorphism is called a *basis* for M. The terminology comes from linear algebra, since a basis of M generates ("spans") it by definition, and is *linearly independent*, i.e. if $x_1, ..., x_n \in X$ are distinct and $\sum r_i x_i = 0$, then $r_i = 0$ for all i. Note that if M is an R-module, and $X \subset M$ is a linearly independent generating (=spanning) set, then $RX \longrightarrow M$ is an isomorphism.

The obvious question that arises is: If M is free, do any two bases for it have the same cardinality? It turns out that the answer is "no" in general. But for most of the rings we consider, the cardinality of a basis is well-defined. To prove this we first consider the case $R = D$ is a division ring, where it turns out that *every* D-module is free.

Theorem 3.2 *Let D be a division ring. Then*
Every D-module M is free. In fact
a) every linearly independent subset of M is contained in a basis.
b) every spanning set (=generating set) contains a basis.

Proof: First note that (a) does imply the theorem, since the empty set is a linearly independent set. Thus M has an D-basis X and is isomorphic to the free module DX.

a) Let L be a linearly independent set, and consider the partially ordered set $\mathcal{S}$ of linearly independent subsets of M that contain L. It is nonempty, since it contains L. Every totally ordered subset $X \subset \mathcal{S}$ has an upper bound B, namely $B = \cup_{A \in X} A$ (note that to check linear independence one need only consider one finite subset at a time, and any finite subset of B lies in some A). So by Zorn's lemma, $\mathcal{S}$ has a maximal element S. I claim that S spans M. For by maximality, for any $v \in M$ we have that $S \cup \{v\}$ is a linearly dependent set, and since S is linearly independent there must be a linear relation of the form

$$c_0 v + \sum_{i=1}^{n} c_i w_i = 0$$

with $w_i \in S$ and $c_0 \neq 0$. *Since D is a division ring*, c_0 is a unit and we can solve the displayed equation for v. Thus v is in the span of S, completing the proof.

b) Let X be a spanning set. A similar Zorn's lemma argument shows that X contains a maximal linearly independent subset S. It follows easily that S is also a spanning set, hence a basis; details are left to the reader.

Theorem 3.3 *Let M be a module over a division ring D. Then any two bases of M have the same cardinality.*

A complete proof of the theorem can be found e.g. in [Hungerford]. Our immediate interest is in the case that M has a finite basis, which can be proved directly by fiddling around with linear combinations of elements. In the interests of efficiency, I'll omit this argument because I'll give a simpler, more conceptual proof later (after discussing Jordan-Holder filtrations).

For the next theorem we need the following observation: If M is an R-module and I is a 2-sided ideal of R, let IM denote the abelian subgroup of M generated by all elements of the form rx with $r \in I$ and $x \in M$. Then IM is an R-submodule and we can form the quotient module M/IM. Moreover this R-module structure factors through an R/I-module structure, namely $(r + I) \cdot (x + IM) = rx + IM$. For example when $R = \mathbb{Z}$ and $I = (n)$, this just says the abelian group M/nM is killed by n and so is a $\mathbb{Z}/n$-module. In general, we need a 2-sided ideal so that R/I is a ring.

By direct check, this construction commutes with direct sums: $(\oplus M_\alpha)/I(\oplus M_\alpha) \cong \oplus(M_\alpha/IM_\alpha)$.

Theorem 3.4 *Let R be a ring, and suppose R has a 2-sided ideal I such that $D := R/I$ is a division ring. Then if M is a free R-module, any two bases of M have the same cardinality.*

Proof: Let X be a basis for M, so that $M = \oplus_{x \in X} R_x$. Then $M/IM = \oplus_{x \in X} D_x$. Hence the cardinality of X is the same as the cardinality of a D-basis for M/IM, and hence is independent of the choice of X by the previous theorem.

Here is the case we will most frequently use:

Corollary 3.5 *Let R be a commutative ring and let M be a free R-module. Then any two bases for M have the same cardinality.*

Proof: Recall that any ring has at least one maximal left ideal I. In a commutative ring I is automatically a 2-sided ideal, and R/I is a field. Now apply the theorem.

When it is well-defined, the cardinality of a basis is called the *rank* of the free module.

Remark. Curiously, the cardinality of a basis is always well-defined when there is an infinite basis (see [Hungerford]). On the other hand when $R = End_F V$ for F a field, V a vector space of countably infinite dimension, one can show $R \cong R \oplus R$ as left R-modules (exercise, if you're curious; it boils down to the usual "Cantor's hotel" trick of splitting the natural numbers into evens and odds). So there is no well-defined rank in this case.

4 Simple modules

A nonzero R-module M is *simple* or *irreducible* (the terms are synonymous) if M has no proper nonzero submodules. A simple module is necessarily cyclic; indeed more is true:

Proposition 4.1 *The following are equivalent:*
a) M is simple;
b) every nonzero element of M generates M;
c) $M \cong R/I$ where I is a maximal left ideal.

Proof: Exercise.

Examples. 1. If R is a field, a module (=vector space) over R is simple if and only if it has dimension 1. In particular, up to isomorphism there is only one simple module. The same assertion holds for any division ring.

2. If R is an F-algebra, any module of F-dimension 1 is simple. For example, any 1-dimensional representation of G over F is a simple FG-module.

3. If R is a principal ideal domain that is not a field, the simple modules are those of the form $R/(p)$ with p irreducible. This is by part (c) of the preceding proposition, since in a principal ideal domain the maximal ideals are precisely those of the form (p), with p irreducible. For example if $f(x) \in F[x]$ is an irreducible polynomial, $F[x]/(f(x))$ is a simple $F[x]$-module.

4. If F is a field, F^n is simple as a module over $M_n F$. (Exercise.)

The next result is known as Schur's lemma. It is completely trivial, yet amazingly powerful.

Proposition 4.2 *If M, N are simple modules and $\phi : M \longrightarrow N$ is a homomorphism, then ϕ is either an isomorphism or zero. Hence:*

a) If M, N are nonisomorphic, $Hom_R(M, N) = 0$.

b) If M is simple, $End_R(M, M)$ is a division ring.

Proof: The kernel of ϕ is either all of M, or zero. In the former case $\phi = 0$ and in the latter ϕ is injective, in which case $Im\,\phi$ must be all of N and ϕ is an isomorphism. This proves the first statement, and (a) follows immediately. Part (b) also follows immediately, since every nonzero homomorphism $M \longrightarrow M$ is invertible.

Remark: Later we will see some interesting examples where the division ring in (b) is the ring of quaternions $\mathbb{H}$. For the time being, however, our examples will have $End_R M$ commutative, i.e. a field. For example, if R is commutative and $M = R/I$ for a maximal ideal I, then R/I is itself a field, and $End_R(R/I) \cong R/I$.

Simple modules serve as building blocks for more general modules. A *Jordan-Holder filtration* of a module is a finite filtration

$$0 = M_0 \subset M_1 \subset ... \subset M_n = M$$

such that each quotient M_i/M_{i-1} is a simple module, $1 \leq i \leq n$. These quotients are called *composition factors*. Note that the zero module is not a simple module in our definition, but by convention we regard the zero module as having the "empty" Jordan-Holder filtration $0 = M_0 = M$.

Proposition 4.3 *The class of modules admitting a Jordan-Holder filtration is closed under taking submodules, quotient modules, finite direct sums, and extensions.*

Proof: Exercise. "Closed under extensions" means that if the two end terms of a short exact sequence admit Jordan-Holder filtrations, so does the middle term.

There are many interesting classes of modules that admit Jordan-Holder filtrations:

Examples. 1. Suppose F is a field and R is an F-algebra. Then every finite dimensional R-module M admits a Jordan-Holder filtration: Any nonzero submodule M_1 of minimal dimension is simple, and by induction on dimension we can assume M/M_1 has a Jordan-Holder filtration $N_2 \subset ... \subset N_n = M/M_1$. Let M_i be the inverse image of N_i in M; then $M_1 \subset M_2 ... \subset M_n = M$ is the desired filtration.

2. Any finite R-module M has a Jordan-Holder filtration. The proof is exactly as in the previous example, using cardinality in place of dimension. For example, any finite abelian group has a finite filtration with quotients of the form $\mathbb{Z}/p$, p ranging over some finite set of primes.

3. Let R be a principal ideal domain. Then any finitely-generated torsion module M admits a Jordan-Holder filtration. Here we are *not* assuming the classification theorem for such modules, which we have not yet proved. So we proceed as follows: First, the assertion is true for cyclic modules $R/(a)$ $(a \neq 0)$. This is easily seen by induction on the number

of factors in a prime factorization of a. So by Proposition 4.3, the assertion is true for any quotient of a direct sum of cyclic modules. But any finitely-generated torsion module is such a quotient: for if $x_1, ..., x_n$ is a generating set, then the evident homomorphism $\oplus Rx_i \longrightarrow M$ is surjective.

There are also many interesting classes of modules which do not admit a Jordan-Holder filtration. In fact, if M has a Jordan-Holder filtration then M is finitely-generated (an easy exercise). So non-finitely-generated modules have no such filtration. More interesting is the following:

Example: Suppose R is a principal ideal domain that is not a field. Then R is not simple, and any nonzero submodule of R is isomorphic to R (as a module). So R has no simple submodules whatsoever, and hence certainly has no Jordan-Holder filtration .

If M admits a Jordan-Holder filtration, then the composition factors which occur are unique up to ordering, in the following sense:

Proposition 4.4 *Suppose M has two Jordan-Holder filtrations, $M_1 \subset ...M_m$ and $N_1 \subset ... \subset N_n$. Then $m = n$, and up to re-ordering and isomorphism, the two sequences of composition factors $M_1, M_2/M_1, ..., M/M_{m-1}$, $N_1, N_2/N_1, ..., M/N_{n-1}$ are identical.*

Proof: By induction on m. If $m = 1$, then $M = M_1$ itself is a simple module, forcing $n = 1$ and $N_1 = M$. At the inductive step, let k be minimal such that $M_1 \subset N_k$. Let f denote the composite $M_1 \longrightarrow N_k \longrightarrow N_k/N_{k-1}$. Then f is nonzero by the minimality of k, and therefore is an isomorphism by Schur's lemma. Thus $M_1 \cong N_k/N_{k-1}$ and $N_k = N_{k-1} \oplus M_1$. We then define a new Jordan-Holder filtration N_i' of M by

$$N_i' = \begin{cases} M_1 \oplus N_{i-1} & \text{if } i \leq k \\ N_i & \text{if } i > k \end{cases}$$

Note that up to isomorphism, the N' filtration has the same composition factors as the N filtration, up to a cyclic permutation of the factors in the first k positions. We have therefore reduced to the case where the M and N filtrations have the same initial term $M_1 = N_1$. So we have two Jordan-Holder filtrations on the quotient M/M_1, to which the inductive hypothesis can be applied, completing the proof.

As a corollary we see that the length n of a Jordan-Holder filtration is independent of the choice of filtration. Hence we may define the *length* of a Jordan-Holder module by $\ell(M) = n$. This is a useful concept in itself, as it provides a substitute for dimension (in the case of finite dimensional modules over F-algebras) that can be used in inductive arguments. Note also:

Proposition 4.5 *Let $0 \longrightarrow L \longrightarrow M \longrightarrow N \longrightarrow 0$ be a short exact sequence of R-modules. Then M has a Jordan-Holder filtration if and only if L, N have Jordan-Holder filtrations, in which case*

$$\ell(M) = \ell(L) + \ell(N).$$

The proof is an easy exercise, following a pattern we've already used for filtrations of groups. As a consequence we obtain a variant of dimension-counting:

Proposition 4.6 *Suppose $\phi : M \longrightarrow N$ is a homomorphism of R-modules of finite length, and $\ell(M) = \ell(N)$. Then ϕ is an isomorphism $\Leftrightarrow$ ϕ is injective $\Leftrightarrow$ if and only if ϕ is surjective.*

Proof: Suppose ϕ is surjective. Then since $\ell(M) = \ell(Ker\,\phi) + \ell(N)$, we have $\ell(Ker\phi) = 0$ and hence $Ker\,\phi = 0$. The case ϕ injective is similar.

Finally, here is the promised proof that the dimension of a finite-dimensional vector space over a division ring is well-defined. To be precise, what we will prove is that if V is a vector space (=module) over a division ring D, and V has a finite basis $v_1, ..., v_n$, then every basis of V has cardinality n. Note first that every cyclic D-module M is simple: If x generates M, then for all $r \in D$, rx also generates M since $x = r^{-1}rx$. Filtering V by the ordered basis v_i, i.e. $0 \subset \langle v_1 \rangle \subset \langle v_1, v_2 \rangle \subset ...$, the quotients of the filtration are cyclic (and nonzero) and hence simple. So $n = \ell(V)$ and hence all finite bases have the same cardinality. In fact what this argument shows is that no linearly independent subset of V has more than n elements (otherwise $\ell(V) > n$), thereby ruling out infinite bases as well.

For the complete result, see the next section.

5 Appendix: Invariance of infinite rank

The goal of this section is to prove:

Theorem 5.1 *Let R be any ring, and let M be a free R-module admitting an infinite basis. Then any two bases of M have the same cardinality.*

Note this completely settles the division ring case, since we just proved the analogous statement for modules admitting a finite basis.

For the proof of the theorem, we will need (unsurprisingly) the Schroder-Bernstein theorem, which says that if X and Y are sets such that $|X| \leq |Y|$ and $|Y| \leq |X|$, then $|X| = |Y|$. The proof of the Schroder-Bernstein theorem is more interesting than you might think, and can be found in one of our standard undergraduate texts (for Mathematical Reasoning, Math 300): *Mathematical Thinking*, by D'Angelo and West. Now, on to the proof of the theorem.

Let X, Y be bases for M, with Y infinite. For every $x \in X$ there is a finite subset $Y_x \subset Y$ such that $x \in Span(Y_x)$. Set

$$Y' = \cup_{x \in X} Y_x.$$

Then $X \subset Span(Y')$ and hence $Span(Y') = M$. But $Y' \subset Y$ and Y is a basis, so no proper subset of Y can span M. Hence $Y' = Y$. From this we conclude two things: First, X is infinite (otherwise Y' would be finite). Second, $|Y| \leq |X|$. This is because $|Y| = |Y'|$ and whenever one has a collection of finite sets A_α indexed by an infinite set J ($J = X$, $A_\alpha = Y_x$ here), $|\cup_{\alpha \in J} A_\alpha| \leq |J|$. Since we now know that X is infinite, the same argument shows $|X| \leq |Y|$ and we are done by Schroder-Bernstein.

6 Exercises

1. *An example of a ring not isomorphic to its opposite.* Consider the ring $\mathfrak{b}_2\mathbb{R}$ of upper triangular 2×2 matrices over $\mathbb{R}$:

$$\begin{pmatrix} a & b \\ 0 & c \end{pmatrix}$$

Let $R \subset \mathfrak{b}_2\mathbb{R}$ denote the set of matrices with $a \in \mathbb{Q}$. Check that R is a subring (you don't need to write it up). Then use the following method to show R is not isomorphic to R^{op}:

1. Show that every right ideal of R is finitely-generated (as right R-module), whereas there is a left ideal that is not finitely-generated.

2. Explain briefly why (1) implies the result.

Remark. Our only interest in this ring is as a counterexample. We will recycle it later as a counterexample to various overly optimistic assertions. On the other hand, the exercise is also intended as practice in working with left/right ideals, and with generators of an ideal (a special case of generators of a module).

2. Show that the only non-trivial group with no non-identity automorphisms is C_2. (Bases of vector spaces will enter at one step, hence, contrary to appearances, the exercise does belong in this section!)

3. Let $0 \longrightarrow L \longrightarrow M \longrightarrow N \longrightarrow 0$ be a short exact sequence of R-modules. Show that if N is a free module, then the short exact sequence splits.

Comment. Criteria guaranteeing that certain short exact sequences split are very helpful, since then we know M completely as $M \cong L \oplus N$. The fact that short exact sequences need not split in general marks the beginning of an entire subject known as "homological algebra". We will make occasional excursions into homological algebra throughout the year.

4. Let F be a field, and regard F^n as a left FS_n-module via permutation of coordinates (this example is discussed in the notes above). Thinking of FS_n-modules as representations of S_n, this corresponds to the inclusion $S_n \subset GL_nF$ as the permutation matrices. Assume $n \geq 2$.

Define $\epsilon : F^n \longrightarrow F$ by $\epsilon(a_1, ..., a_n) = \sum a_i$, and note that ϵ is an FS_n-module map, where S_n acts trivially on F. Let $V = Ker\, \epsilon$, so there is a short exact sequence of FS_n-modules

$$0 \longrightarrow V \longrightarrow F^n \xrightarrow{\epsilon} F \longrightarrow 0.$$

a) Show that this sequence splits if and only if *char* F doesn't divide n.

b) For $n \geq 3$, show that V is a simple module if and only if *char* F doesn't divide n. (For $n = 2$ it is automatically simple, by virtue of being 1-dimensional.)

5. Regard F^n as a left M_nF-module as usual, thinking of elements of F^n as column vectors. Prove the following:

a) F^n is a simple module.

b) As a left module over itself, $M_nF = \oplus_{i=1}^{n} F^n$ (the direct sum of n copies of F^n).

c) Up to isomorphism, F^n is the only simple M_nF-module.