19 Modules over a Ring

The motivation for modules is that we are trying to tell a story where rings are the protagonist, and for a story to be interesting, the protagonist must act. When we find a way for a ring to act, we get the definition of a module.

Definition 19.1
Let R be a ring. A **module** M over R is an abelian group, together with an **action map** $R \times M \to M$ (written as $(r, m) \mapsto r(m)$ or rm), subject to the following axioms:

- $1_R(m) = m$ for all m im M;
- $r_1(r_2(m)) = (r_1r_2)(m)$ for all $r_1, r_2 \in R$ and $m \in M$;
- Distributivity in both variables: $(r_1 + r_2)m = r_1m + r_2m$ for all $r_1, r_2 \in R$ and $m \in M$, and $r(m_1 + m_2) = rm_1 + rm_2$ for all $r \in R$ and $m_1, m_2 \in M$.

The first two axioms are very similar to the definition of a group action on a set. So a ring to a module is like a group G to a G-set (a set with an action by G). It's not exactly the same because in a ring we have two operations instead of one, but they're a similar flavor.

19.1 Examples

Example 19.2
If $R = F$ is a field, then a module is the same as a vector space.

The axioms here are exactly the same. The textbook emphasizes this heavily, and this analogy can get you some mileage; but for general rings, things become more complicated.

Note 19.3
The definition also applies to a *noncommutative* ring R, in the same way — our definition does not reference commutativity. Then we have some familiar examples of modules over noncommutative rings: for example, given any field F we can take $R = \mathrm{Mat}_{n\times n}(F)$ and $M = F^n$, since matrices act on column vectors by multiplication. As another example, if $R = \mathbb{C}[G]$ is the group ring, then a R-module is the same as a complex representation of G.

For any ring R, there is a uniquely defined homomorphism $\mathbb{Z} \to R$, where $1 \mapsto 1_R$. On a similar note, every abelian group has a unique structure of a $\mathbb{Z}$-module: we know that 1 (in $\mathbb{Z}$) must map $m \mapsto m$, so then by distributivity, $n = 1 + 1 + \cdots + 1$ must map

$$v \mapsto \underbrace{v + v + \cdots + v}_{n}.$$

Similarly, $-n$ must map v to $-(v + v + \cdots + v)$. So a $\mathbb{Z}$-module is the same as an abelian group.

Example 19.4
What is a module over $R = \mathbb{C}[x]$?

Proof. First, a $\mathbb{C}[x]$-module is a $\mathbb{C}$-vector space V by looking at the action of constant polynomials (which are just scalars). But then we also need to see what x does. We know x must act by a linear map $A : V \to V$, where $xv = Av$. There are no constraints on this map, and this defines the action of every other polynomial: so a R-module is a vector space V, together with a linear map $A : V \to V$. Explicitly, the action of a general polynomial $P(x) = a_nx^n + \cdots + a_0$ is given by

$$Pv = a_0v + a_1Av + \cdots + a_nA^nv.$$

Note that the vector space may or may not be finite-dimensional; if it is, then we end up in a situation studied in linear algebra, where we have a vector space and a linear operator. $\square$

Example 19.5
What is a module over $R = \mathbb{Z}/n\mathbb{Z}$?

Proof. The main point is that if R/I is a quotient of R, then every R/I-module is also a R-module, where we define $r(m)$ to be $\overline{r}(m)$ (here $\overline{r}$ denotes r mod I). Meanwhile, in order to go backwards, I must act by 0. So a R/I module is the same as a R-module where every element of I acts in a trivial way (meaning that $rv = 0$ for all $r \in I$ and $v \in M$).

So in this case, a $\mathbb{Z}/n\mathbb{Z}$-module is the same as an abelian group where the order of every element divides n — meaning $na = 0$ for all a in the group.

Then more concretely, for every m (where we use $\overline{m}$ to denote m mod n), we can write

$$\overline{m}v = \underbrace{v + v + \cdots + v}_{m}.$$

In order for this to be well-defined, the sum should not depend on the choice of representative for the residue; but this is guaranteed by the condition $na = 0$. (This is the same reasoning as in the first paragraph, for this specific example.) □

For any ring R, there is a simple example of a module:

Definition 19.6
The **free module** over R is $M = R$ itself, where the action is multiplication (meaning that $r(x) = rx$).

This is parallel to the observation that a group G acts on itself by left multiplication.

19.2 Submodules

Definition 19.7
Given a module M, a **submodule** $N \subset M$ is an abelian subgroup which is invariant under the R-action — meaning $rx \in N$ for all $x \in N$ and $r \in R$.

If $N \subset M$ is a submodule, we can define their **quotient** M/N, where we take the quotient in the sense of abelian groups. This quotient of abelian groups carries a module structure as well, given by the obvious rule $r\overline{m} = \overline{rm}$ (where $\overline{m}$ denotes m mod N). This is well-defined because N is a submodule — if $m_1 - m_2$ is in N, then $rm_1 - rm_2 = r(m_1 - r_2)$ is in N as well.

Then the homomorphism theorem and correspondence theorem work in the exact same way as in abelian groups. (For rings and ideals, we saw they work in a similar way; but here the parallel is closer.)

Example 19.8
What are the submodules of the free module $M = R$?

Proof. The answer is exactly the ideals of R — we're looking for abelian subgroups of R which are invariant under multiplication by all terms in R, and by definition these are ideals. □

We'll later see how to understand *any* module by looking at generators and relations — this turns out to be easier than the corresponding problem for a group. But first we'll look at another example of a module, which will be useful for developing that theory.

Definition 19.9
Given two modules M and N, their **direct sum** is

$$M \oplus N = \{(m, n) \mid m \in M, n \in N\}$$

with the action

$$r(m, n) = (rm, rn).$$

Note 19.10
The direct sum is the same as the product $M \times N$. This is true for any *finite* sum — we have

$$M_1 \oplus \cdots \oplus M_n = M_1 \times \cdots \times M_n.$$

But this isn't true for *infinite* sums and products.

Definition 19.11
The **free module** of **rank** n is

$$R^n = \underbrace{R \oplus R \oplus \cdots \oplus R}_{n}.$$

In the case where $R = F$ is a field, the free module of rank n is exactly F^n, the standard n-dimensional vector space.

19.3 Homomorphisms

In linear algebra, we work with matrices in order to understand linear maps. Matrices are also relevant here — the terms are different, but the concept is very similar.

Definition 19.12
A **homomorphism** from a module M to a module N is a homomorphism of abelian groups $\varphi : M \to N$, which is compatible with the R-action — meaning $\varphi(rm) = r\varphi(m)$ for all $r \in R$ and $m \in M$.

In vector spaces, this is the same as a linear map.

We'll use $\mathrm{Hom}_R(M, N)$ to denote the set of all homomorphisms $M \to N$. Note that homomorphisms can be added and rescaled, in the same way as linear maps: $(\varphi_1 + \varphi_2)(m) = \varphi_1(m) + \varphi_2(m)$, and $(r\varphi)(m) = r\varphi(m)$. So then $\mathrm{Hom}_R(M, N)$ is itself a R-module.

Understanding homomorphisms in general may be hard, but it's easy to understand homomorphisms from a free module. Given a homomorphism $\varphi \in \mathrm{Hom}_R(R, M)$ for any module M, we can let $m = \varphi(1_R)$. Then this determines the entire homomorphism — for any $r \in R$, we have

$$\varphi(r) = \varphi(r \cdot 1_R) = r \cdot \varphi(1_R) = rm.$$

So a homomorphism is determined by $m = \varphi(1_R)$, and there are no restrictions on m — this is why R is called a free module. This means $\mathrm{Hom}_R(R, M)$ is isomorphic to M: more explicitly, the bijection is given by mapping $\varphi \in \mathrm{Hom}_R(R, M)$ to $m_\varphi = \varphi(1)$, and $m \in M$ to the homomorphism $\varphi_m : r \mapsto rm$.

Similarly, $\mathrm{Hom}_R(R^n, M)$ is equally easy to understand. Now R^n is generated by the elements 1_i which have a 1 in their ith place, and 0's everywhere else (so $1_i = (0, \dots, 0, 1, 0, \dots, 0)$, where the 1 is in the ith place). So $\mathrm{Hom}_R(R^n, M)$ is isomorphic to M^n, where the bijection sends $\varphi \in \mathrm{Hom}_R(R^n, M)$ to the element $(\varphi(1_1), \varphi(1_2), \dots, \varphi(1_n))$, and $(m_1, \dots, m_n) \in M$ to the homomorphism $\varphi(x_1, \dots, x_n) = \sum x_i m_i$.

In particular, we have $\mathrm{Hom}_R(R^n, R^m) = (R^m)^n = \mathrm{Mat}_{m\times n}(R)$ — we can write homomorphisms in the way we're used to in linear algebra, where $A \in \mathrm{Mat}_{m\times n}(R)$ sends $(x_1, \dots, x_n)^t$ to $A(x_1, \dots, x_n)^t$. So as long as we work with free modules and homomorphisms, to a large extent we can operate as if we're doing linear algebra. But in linear algebra, there's various characterizations of nondegenerate matrices that no longer hold here —

for instance, a linear operator that is injective (meaning it has zero kernel) is also surjective, but that's not true for general modules.

19.4 Generators and Relations

Definition 19.13
A collection of elements $m_1, \ldots, m_n \in M$ forms a system of **generators** if every $x \in M$ can be expressed as $\sum r_i m_i$ for $r_i \in R$.

So in other words, $\varphi_{m_1,\ldots,m_n} : R^n \to M$ is onto. If such a finite set exists, we say that M is *finitely generated.* Many modules we're interested in are in fact finitely generated.

If this map is also one-to-one, then it's an isomorphism, and M is free. But usually this won't happen, and we still want to describe M. To do this, we can look at $K = \ker(\varphi)$, which is a submodule in R^n. If K is itself finitely generated, then we can choose a set of generators for K, and get a somewhat explicit description of M — we can fix a system of k generators for K, and obtain another homomorphism $\psi : R^k \to K$. Since K sits inside R^n, we can think of ψ instead as a homomorphism $\psi : R^k \to R^n$, with image K. But such a homomorphism corresponds to a matrix $A \in \text{Mat}_{k\times n}$, and we have $M = R^n/AR^n$.

Definition 19.14
If we can find a finite set of generators for M such that the kernel is also finitely generated, then M is called **finitely presented**.

We'll see that for many rings, the finitely presented modules are a very large class of modules — in fact, for many rings, any finitely generated module is finitely presented. We'll also see how to make this description of a module very explicit when the ring is a Euclidean domain. In particular, taking the Euclidean domain to be $\mathbb{Z}$ will give us a classification of all finitely generated abelian groups, and taking the Euclidean domain to be $F[x]$ will give us Jordan normal form!