

EXTERIOR POWERS

KEITH CONRAD

1. INTRODUCTION

Let R be a commutative ring. Unless indicated otherwise, all modules are R -modules and all tensor products are taken over R , so we abbreviate \otimes_R to \otimes . A bilinear function out of $M_1 \times M_2$ turns into a linear function out of the tensor product $M_1 \otimes M_2$. In a similar way, a multilinear function out of $M_1 \times \cdots \times M_k$ turns into a linear function out of the k -fold tensor product $M_1 \otimes \cdots \otimes M_k$. We will concern ourselves with the case when the component modules are all the same: $M_1 = \cdots = M_k = M$. The tensor power $M^{\otimes k}$ universally linearizes all the multilinear functions on M^k .

A function on M^k can have its variables permuted to give new functions on M^k . When looking at permutations of the variables, two important types of functions on M^k occur: symmetric and alternating. These will be defined in Section 2. We will introduce in Section 3 the module which universally linearizes the alternating multilinear functions on M^k : the exterior power $\Lambda^k(M)$. It is a certain quotient module of $M^{\otimes k}$. The special case of exterior powers of finite free modules will be examined in Section 4. Exterior powers will be extended from modules to linear maps in Section 5. Applications of exterior powers to determinants are in Section 6 and to linear independence are in Section 7. Section 8 will introduce a product $\Lambda^k(M) \times \Lambda^\ell(N) \rightarrow \Lambda^{k+\ell}(M)$. Finally, in Section 9 we will combine all the exterior powers of a fixed module into a noncommutative ring called the exterior algebra of the module.

The exterior power construction is important in geometry, where it provides the language for discussing differential forms on manifolds. (A differential form on a manifold is related to exterior powers of the dual space of the tangent space of a manifold at each of its points.) Exterior powers also arise in representation theory, as one of several ways of creating new representations of a group from a given representation of the group. In linear algebra, exterior powers provide an algebraic mechanism for detecting linear relations among vectors and for studying the “geometry” of the subspaces of a vector space.

2. SYMMETRIC AND ALTERNATING FUNCTIONS

For any function $f: M^k \rightarrow N$ and any $\sigma \in S_k$, we get a new function $M^k \rightarrow N$ by permuting the variables in f according to σ :

$$(m_1, \dots, m_k) \mapsto f(m_{\sigma(1)}, \dots, m_{\sigma(k)}) \in N.$$

(Warning: if we regard this new function on M^k as the effect of σ on f , and write it as $(\sigma \cdot f)(m_1, \dots, m_k)$, then $\sigma_1 \cdot (\sigma_2 \cdot f)$ equals $(\sigma_2 \sigma_1) \cdot f$, *not* $(\sigma_1 \sigma_2) \cdot f$, so we don't have a left action of S_k on the functions $M^k \rightarrow N$ but a right action. We won't be using group actions, so don't worry about this.)

Definition 2.1. We call $f: M^k \rightarrow N$ *symmetric* if

$$f(m_{\sigma(1)}, \dots, m_{\sigma(k)}) = f(m_1, \dots, m_k)$$

for all $\sigma \in S_k$. We call f *skew-symmetric* if

$$(2.1) \quad f(m_{\sigma(1)}, \dots, m_{\sigma(k)}) = (\text{sign } \sigma) f(m_1, \dots, m_k)$$

for all $\sigma \in S_k$. For $k \geq 2$, f is called *alternating* if

$$f(m_1, \dots, m_k) = 0 \text{ whenever } m_i = m_j \text{ for some } i \neq j.$$

Symmetric functions are unchanged by any permutation of the variables, while skew-symmetric functions are unchanged by even permutations and change sign under odd permutations. For example, the value of a skew-symmetric function changes by a sign if we permute any two of the variables. Alternating functions might not seem so intuitive. When R is a field of characteristic 0, like the real numbers, we will see that alternating and skew-symmetric *multilinear* functions are the same thing (Theorem 2.10).

Example 2.2. The function $M_n(R) \times M_n(R) \rightarrow R$ by $(A, B) \mapsto \text{Tr}(AB)$ is symmetric.

Example 2.3. The function $R^2 \times R^2 \rightarrow R$ given by $\left(\begin{pmatrix} a \\ c \end{pmatrix}, \begin{pmatrix} b \\ d \end{pmatrix}\right) \mapsto ad - bc$ is skew-symmetric and alternating.

Example 2.4. The cross product $\mathbf{R}^3 \times \mathbf{R}^3 \rightarrow \mathbf{R}^3$ is skew-symmetric and alternating.

Example 2.5. The function $\mathbf{C} \times \mathbf{C} \rightarrow \mathbf{R}$ given by $(z, w) \mapsto \text{Im}(z\bar{w})$ is skew-symmetric and alternating.

Example 2.6. Let R contain $\mathbf{Z}/2\mathbf{Z}$, so $-1 = 1$ in R . The multiplication map $R \times R \rightarrow R$ is symmetric and skew-symmetric, but *not* alternating.

In Definition 2.1, the variables are indexed in the order from 1 to n . Let's show the properties of symmetry, skew-symmetry, and alternating don't depend on this particular ordering.

Theorem 2.7. Fix a listing of the numbers from 1 to k as i_1, i_2, \dots, i_k . If a function $f: M^k \rightarrow N$ is symmetric then

$$f(m_{\sigma(i_1)}, \dots, m_{\sigma(i_k)}) = f(m_{i_1}, \dots, m_{i_k})$$

for all $\sigma \in S_k$. If f is skew-symmetric then

$$f(m_{\sigma(i_1)}, \dots, m_{\sigma(i_k)}) = (\text{sign } \sigma) f(m_{i_1}, \dots, m_{i_k})$$

for all $\sigma \in S_k$. For $k \geq 2$, if f is alternating then

$$f(m_{i_1}, \dots, m_{i_k}) = 0 \text{ whenever } m_{i_s} = m_{i_t} \text{ for some } i_s \neq i_t.$$

Proof. We will discuss the skew-symmetric case, leaving the other two cases to the reader. Let $\tilde{\sigma} \in S_k$ be the permutation where $\tilde{\sigma}(1) = i_1, \dots, \tilde{\sigma}(k) = i_k$. If f is skew-symmetric,

$$\begin{aligned} f(m_{i_1}, \dots, m_{i_k}) &= f(m_{\tilde{\sigma}(1)}, \dots, m_{\tilde{\sigma}(k)}) \\ &= (\text{sign } \tilde{\sigma}) f(m_1, \dots, m_k), \end{aligned}$$

so for any $\sigma \in S_k$

$$\begin{aligned} f(m_{\sigma(i_1)}, \dots, m_{\sigma(i_k)}) &= f(m_{(\sigma\tilde{\sigma})(1)}, \dots, m_{(\sigma\tilde{\sigma})(k)}) \\ &= \text{sign}(\sigma\tilde{\sigma})f(m_1, \dots, m_k) \\ &= \text{sign}(\sigma)\text{sign}(\tilde{\sigma})f(m_1, \dots, m_k) \\ &= \text{sign}(\sigma)f(m_{i_1}, \dots, m_{i_k}). \end{aligned}$$

□

The next two theorems explain the connection between alternating multilinear functions and skew-symmetric multilinear functions, which is suggested by the above examples.

Theorem 2.8. *For $k \geq 2$, a multilinear function $f: M^k \rightarrow N$ which is alternating is skew-symmetric.*

Proof. We first do the case $k = 2$, because the basic idea is already evident there. When $f: M^2 \rightarrow N$ is alternating, $f(m, m) = 0$ for all m in M . So for all m and m' in M ,

$$f(m + m', m + m') = 0.$$

Expanding by linearity in each component,

$$f(m, m) + f(m, m') + f(m', m) + f(m', m') = 0.$$

The first and last terms are 0, so $f(m, m') = -f(m', m)$. This means f is skew-symmetric.

For the general case when $k \geq 2$, we want to show (2.1) for all $\sigma \in S_k$. Notice first that if f satisfies (2.1) for the permutations σ_1 and σ_2 in S_k then

$$\begin{aligned} f(m_{(\sigma_1\sigma_2)(1)}, \dots, m_{(\sigma_1\sigma_2)(k)}) &= f(m_{\sigma_1(\sigma_2(1))}, \dots, m_{\sigma_1(\sigma_2(k))}) \\ &= (\text{sign } \sigma_1)f(m_{\sigma_2(1)}, \dots, m_{\sigma_2(k)}) \\ &= (\text{sign } \sigma_1)(\text{sign } \sigma_2)f(m_1, \dots, m_k) \\ &= \text{sign}(\sigma_1\sigma_2)f(m_1, \dots, m_k). \end{aligned}$$

Hence to verify (2.1) for all $\sigma \in S_k$ it suffices to verify (2.1) as σ runs over a generating set of S_k . We will check (2.1) when σ runs over the generating set of transpositions

$$\{(i \ i + 1) : 1 \leq i \leq k - 1\}.$$

That is, for $(m_1, \dots, m_k) \in M^k$, we want to show any alternating multilinear function $f: M^k \rightarrow N$ satisfies

$$(2.2) \quad f(\dots, m_{i-1}, m_i, m_{i+1}, m_{i+2}, \dots) = -f(\dots, m_{i-1}, m_{i+1}, m_i, m_{i+2}, \dots),$$

where we only interchange the places of m_i and m_{i+1} .

Fix all components for f except those in positions i and $i + 1$, reducing us to a function of two variables: choose $m_1, \dots, m_{i-1}, m_{i+2}, \dots, m_k \in M$ (an empty condition if $k = 2$) and let $g(x, y) = f(m_1, \dots, m_{i-1}, x, y, m_{i+2}, \dots)$. Then g is bilinear and alternating. Therefore by the $k = 2$ case g is skew-symmetric: $g(x, y) = -g(y, x)$. This implies (2.2), so we're done. □

Corollary 2.9. *For $k \geq 2$, a function $f: M^k \rightarrow N$ is skew-symmetric if and only if it satisfies*

$$f(m_1, \dots, m_{i+1}, m_i, \dots, m_k) = -f(m_1, \dots, m_i, m_{i+1}, \dots, m_k)$$

for $1 \leq i \leq k - 1$, and f is alternating if and only if

$$f(m_1, \dots, m_i, m_{i+1}, \dots, m_k) = 0$$

whenever $m_i = m_{i+1}$ for $1 \leq i \leq k-1$.

Proof. The second paragraph of the proof of Theorem 2.8 applies to all functions, not just multilinear functions, so the condition equivalent to skew-symmetry follows.

If we now suppose f vanishes at any k -tuple with a pair of adjacent equal coordinates, then what we just proved shows f is skew-symmetric. Therefore the value of f at any k -tuple with a pair of equal coordinates is up to sign its value at a k -tuple with a pair of adjacent equal coordinates, and that value is 0 by hypothesis. \square

Example 2.6 shows the converse of Theorem 2.8 can fail: a multilinear function can be skew-symmetric and not alternating. But if 2 is a unit in R (e.g., $R = \mathbf{R}$) then the converse of Theorem 2.8 does hold:

Theorem 2.10. *Let $k \geq 2$. If $2 \in R^\times$ then a multilinear function $f: M^k \rightarrow N$ which is skew-symmetric is alternating.*

Proof. We will show $f(m_1, m_2, \dots, m_k) = 0$ when $m_1 = m_2$. The argument when $m_i = m_j$ for other distinct pairs i and j is the same. By skew-symmetry,

$$f(m_2, m_1, m_3, \dots, m_k) = -f(m_1, m_2, m_3, \dots, m_k).$$

Therefore

$$f(m, m, m_3, \dots, m_k) = -f(m, m, m_3, \dots, m_k),$$

so

$$2f(m, m, m_3, \dots, m_k) = 0.$$

Since 2 is in R^\times , $f(m, m, m_3, \dots, m_k) = 0$. \square

Strictly speaking, the assumption that $2 \in R^\times$ could be weakened to 2 not being a zero divisor in R . (This is stronger than saying R doesn't have characteristic 2, since $R = \mathbf{Z}/(4)$ doesn't have characteristic 2 but the proof doesn't work for such R .)

Over the real and complex numbers the terms "alternating" and "skew-symmetric" can be used interchangeably for multilinear functions. For instance in [5], where only real vector spaces are used, multilinear functions satisfying the skew-symmetric property (2.1) are called alternating. At any point in our discussion where there is a noticeable difference between being alternating and being skew-symmetric, it is just a technicality, so don't worry too much about it.

3. EXTERIOR POWERS OF MODULES

Let M and N be R -modules and $k \geq 1$. When $f: M^k \rightarrow N$ is multilinear and $g: N \rightarrow P$ is linear, the composite $g \circ f$ is multilinear. Moreover, if f is symmetric, skew-symmetric, or alternating then $g \circ f$ has the corresponding property too. So we can create new (symmetric, skew-symmetric, or alternating) multilinear maps from old ones by composing with a linear map.

The k th tensor power $M^k \xrightarrow{\otimes} M^{\otimes k}$, sending (m_1, \dots, m_k) to $m_1 \otimes \dots \otimes m_k$, is a particular example of a multilinear map out of M^k , and every other example comes from this one: given any multilinear map $f: M^k \rightarrow N$ there is a unique linear map $\tilde{f}: M^{\otimes k} \rightarrow N$ whose composite with $M^k \xrightarrow{\otimes} M^{\otimes k}$ is f . That is, there is a unique linear map $\tilde{f}: M^{\otimes k} \rightarrow N$

making the diagram

$$\begin{array}{ccc} M^k & \xrightarrow{f} & N \\ \otimes \downarrow & \nearrow \tilde{f} & \\ M^{\otimes k} & & \end{array}$$

commute, which means $\tilde{f}(m_1 \otimes \cdots \otimes m_k) = f(m_1, \dots, m_k)$.

We now focus our attention on multilinear f which are alternating. When $k \geq 2$ and $f: M^k \rightarrow N$ is alternating, f vanishes on any k -tuple with a pair of equal coordinates, so $\tilde{f}: M^{\otimes k} \rightarrow N$ vanishes on any tensor $m_1 \otimes \cdots \otimes m_k$ with $m_i = m_j$ for some $i \neq j$. Thus the submodule J_k spanned by these special tensors is automatically in the kernel of \tilde{f} . The quotient of $M^{\otimes k}$ by the submodule J_k will be our main object of interest.

Definition 3.1. For an R -module M and an integer $k \geq 2$, the k th exterior power of M , denoted $\Lambda^k(M)$, is the R -module $M^{\otimes k}/J_k$ where J_k is the submodule of $M^{\otimes k}$ spanned by all $m_1 \otimes \cdots \otimes m_k$ with $m_i = m_j$ for some $i \neq j$. For any $m_1, \dots, m_k \in M$, the coset of $m_1 \otimes \cdots \otimes m_k$ in $\Lambda^k(M)$ is denoted $m_1 \wedge \cdots \wedge m_k$. For completeness, set $\Lambda^0(M) = R$ and $\Lambda^1(M) = M$ (so $J_0 = J_1 = \{0\}$).

We could write $\Lambda_R^k(M)$ to place the ring R in our notation. But since we will never be changing the ring, we suppress this extra decoration. The general element of $\Lambda^k(M)$ will be denoted ω or η (as we write t for a general tensor). Since $M^{\otimes k}$ is spanned by the tensors $m_1 \otimes \cdots \otimes m_k$, the quotient module $M^{\otimes k}/J_k = \Lambda^k(M)$ is spanned by their images $m_1 \wedge \cdots \wedge m_k$. That is, any $\omega \in \Lambda^k(M)$ is a finite R -linear combination

$$\omega = \sum r_{i_1, \dots, i_k} m_{i_1} \wedge \cdots \wedge m_{i_k},$$

where the coefficients r_{i_1, \dots, i_k} are in R and the m_i 's are in M .

We call $m_1 \wedge m_2 \wedge \cdots \wedge m_k$ an *elementary wedge product* and read it as “ m_1 wedge $m_2 \dots$ wedge m_k .” Another name for elementary wedge products is decomposable elements. (More synonyms: simple, pure, monomial). Since $r(m_1 \wedge m_2 \wedge \cdots \wedge m_k) = (rm_1) \wedge m_2 \wedge \cdots \wedge m_k$, every element of $\Lambda^k(M)$ is a sum (not just a linear combination) of elementary wedge products. A linear – or even additive – map out of $\Lambda^k(M)$ is completely determined by its values on elementary wedge products because they additively span $\Lambda^k(M)$. More general wedge products will be met in Section 8.

The modules $\Lambda^k(M)$ were introduced by Grassmann (for $M = \mathbf{R}^n$), who called expressions like $m_1 \wedge m_2$ outer products. Now we use the label “exterior” instead. Perhaps it would be better to call $\Lambda^k(M)$ an alternating power instead of an exterior power, but it’s too late to change the terminology.

Example 3.2. Suppose M is spanned by the two elements x and y : $M = Rx + Ry$. (This doesn’t mean x and y are a basis over R , e.g., $R = \mathbf{Z}[\sqrt{-5}]$ and $M = (2, 1 + \sqrt{-5})$.) We will show $\Lambda^2(M)$ is spanned by the single element $x \wedge y$. The tensor square $M^{\otimes 2}$ is spanned by all terms $m \otimes m'$ where m and m' are in M . Write $m = ax + by$ and $m' = cx + dy$. Then in $M^{\otimes 2}$

$$\begin{aligned} m \otimes m' &= (ax + by) \otimes (cx + dy) \\ &= ac(x \otimes x) + ad(x \otimes y) + bc(y \otimes x) + bd(y \otimes y). \end{aligned}$$

The tensors $x \otimes x$ and $y \otimes y$ are in J_2 , so in $\Lambda^2(M)$ both $x \wedge x$ and $y \wedge y$ vanish. Therefore

$$m \wedge m' = ad(x \wedge y) + bc(y \wedge x).$$

Moreover, the tensor $(x + y) \otimes (x + y)$ is in J_2 , so

$$x \otimes y + y \otimes x = (x + y) \otimes (y + x) - x \otimes x - y \otimes y \in J_2.$$

Therefore in $\Lambda^2(M)$ we have $x \wedge y + y \wedge x = 0$, so

$$m \wedge m' = ad(x \wedge y) + bc(-x \wedge y) = (ad - bc)(x \wedge y),$$

which means $\Lambda^2(M)$ is spanned by the single element $x \wedge y$. It could happen that $x \wedge y = 0$ (so $\Lambda^2(M)$ could be zero), or even if $x \wedge y \neq 0$ it could happen that $r(x \wedge y) = 0$ for some nonzero $r \in R$. It all depends on the nature of the R -linear relations between x and y in M .

By comparison to $\Lambda^2(M)$, $M^{\otimes 2}$ is spanned by $x \otimes x$, $x \otimes y$, $y \otimes x$, and $y \otimes y$, and without further information we have no reason to collapse this spanning set (usually $x \otimes y \neq y \otimes x$, for instance). So when M has a 2-element spanning set, a spanning set for $M^{\otimes 2}$ is typically larger than 2 while a spanning set for $\Lambda^2(M)$ is definitely smaller.

For $k \geq 2$, the standard map $M^k \xrightarrow{\otimes} M^{\otimes k}$ is multilinear and the reduction map $M^{\otimes k} \rightarrow M^{\otimes k}/J_k = \Lambda^k(M)$ is linear, so the composite map $\wedge: M^k \xrightarrow{\otimes} M^{\otimes k} \rightarrow \Lambda^k(M)$ is multilinear. That is, the function

$$(3.1) \quad (m_1, \dots, m_k) \mapsto m_1 \wedge \cdots \wedge m_k$$

from M^k to $\Lambda^k(M)$ is multilinear in the m_i 's. For example,

$$m_1 \wedge \cdots \wedge cm_i \wedge \cdots \wedge m_k = c(m_1 \wedge \cdots \wedge m_i \wedge \cdots \wedge m_k), \quad m_1 \wedge \cdots \wedge 0 \wedge \cdots \wedge m_k = 0.$$

Working in $\Lambda^k(M)$ forces $m_1 \wedge \cdots \wedge m_k = 0$ if $m_i = m_j$ for some $i \neq j$. (Think about what working modulo J_k means for the tensors $m_1 \otimes \cdots \otimes m_k$ when $m_i = m_j$ for some $i \neq j$.) Therefore (3.1) is an *example* of an alternating multilinear map out of M^k . Now we show it is a universal example: all others pass through it using linear maps out of $\Lambda^k(M)$.

Theorem 3.3. *Let M be an R -module and $k \geq 2$. For any R -module N and alternating multilinear map $f: M^k \rightarrow N$, there is a unique linear map $\tilde{f}: \Lambda^k(M) \rightarrow N$ such that the diagram*

$$\begin{array}{ccc} M^k & \xrightarrow{f} & N \\ \wedge \downarrow & \nearrow \tilde{f} & \\ \Lambda^k(M) & & \end{array}$$

commutes, i.e., $\tilde{f}(m_1 \wedge \cdots \wedge m_k) = f(m_1, \dots, m_k)$.

This theorem makes no sense when $k = 0$ or $k = 1$ since there are no alternating multilinear functions in those cases.

Proof. Since f is multilinear, it induces a linear map $M^{\otimes k} \rightarrow N$ whose behavior on elementary tensors is

$$m_1 \otimes \cdots \otimes m_k \mapsto f(m_1, \dots, m_k).$$

$$\begin{array}{ccc} M^k & \xrightarrow{f} & N \\ \otimes \downarrow & \nearrow & \\ M^{\otimes k} & & \end{array}$$

(We are going to avoid giving this map $M^{\otimes k} \rightarrow N$ a specific notation, since it is just an intermediate device in this proof.) Because f is alternating, it vanishes at any k -tuple (m_1, \dots, m_k) where $m_i = m_j$ for some $i \neq j$. Thus the linear map which f induces from $M^{\otimes k}$ to N vanishes at any elementary tensor $m_1 \otimes \dots \otimes m_k$ where $m_i = m_j$ for some $i \neq j$. Hence this linear map out of $M^{\otimes k}$ vanishes on the submodule J_k of $M^{\otimes k}$, so we get an induced linear map \tilde{f} out of $M^{\otimes k}/J_k = \Lambda^k(M)$. Specifically, $\tilde{f}: \Lambda^k(M) \rightarrow N$ is given by

$$m_1 \wedge \dots \wedge m_k \mapsto f(m_1, \dots, m_k).$$

$$\begin{array}{ccc} M^k & \xrightarrow{f} & N \\ \downarrow & \nearrow & \\ M^{\otimes k} & & \\ \downarrow & \nearrow \tilde{f} & \\ \Lambda^k(M) & & \end{array}$$

Since the elements $m_1 \wedge \dots \wedge m_k$ span $\Lambda^k(M)$, a linear map out of $\Lambda^k(M)$ is uniquely determined by its effect on these elements. Thus, having constructed a linear map out of $\Lambda^k(M)$ whose effect on any $m_1 \wedge \dots \wedge m_k$ is the same as the effect of f on (m_1, \dots, m_k) , it is the unique such linear map. \square

We will call the particular alternating multilinear map $M^k \xrightarrow{\wedge} \Lambda^k(M)$ given by

$$(m_1, \dots, m_k) \mapsto m_1 \wedge \dots \wedge m_k$$

the *canonical map*.

Remark 3.4. We could have constructed $\Lambda^k(M)$ as the quotient of a huge free module on the set M^k , bypassing the use of $M^{\otimes k}$. Since the canonical map $M^k \xrightarrow{\wedge} \Lambda^k(M)$ is multilinear, we get a linear map $M^{\otimes k} \rightarrow \Lambda^k(M)$ and can recover $\Lambda^k(M)$ as a quotient of $M^{\otimes k}$ anyway.

Corollary 3.5. *Suppose Λ is an R -module and there is an alternating multilinear map $\alpha: M^k \rightarrow \Lambda$ with the same universal mapping property as the canonical map $M^k \xrightarrow{\wedge} \Lambda^k(M)$: for every R -module N and alternating multilinear map $f: M^k \rightarrow N$ there is a unique linear map $\Lambda \rightarrow N$ making the diagram*

$$\begin{array}{ccc} M^k & \xrightarrow{f} & N \\ \alpha \downarrow & \nearrow & \\ \Lambda & & \end{array}$$

commute. Then there is a unique R -linear map $\beta: \Lambda \rightarrow \Lambda^k(M)$ such that the diagram

$$\begin{array}{ccc} & M^k & \\ \alpha \swarrow & & \searrow \wedge \\ \Lambda & \xrightarrow{\beta} & \Lambda^k(M) \end{array}$$

commutes, and β is an isomorphism.

Proof. This is the usual argument that an object equipped with a map satisfying a universal mapping property is determined up to a unique isomorphism: set two such objects and maps against each other to get maps between the objects in both directions whose composites in both orders have to be the identity maps on the two objects by the usual argument. \square

Since the canonical map $M^k \xrightarrow{\wedge} \Lambda^k(M)$ is alternating multilinear, by Theorem 2.8 it is skew-symmetric:

$$(3.2) \quad m_{\sigma(1)} \wedge \cdots \wedge m_{\sigma(k)} = (\text{sign } \sigma) m_1 \wedge \cdots \wedge m_k$$

for every $\sigma \in S_k$. In particular, an elementary wedge product $m_1 \wedge \cdots \wedge m_k$ in $\Lambda^k(M)$ is determined up to an overall sign by the terms m_i appearing in it (e.g., $m \wedge m' \wedge m'' = -m' \wedge m \wedge m''$).

Example 3.6. Returning to Example 3.2 and working directly in $\Lambda^2(M)$ from the start, we have

$$(3.3) \quad (ax + by) \wedge (cx + dy) = ac(x \wedge x) + ad(x \wedge y) + bc(y \wedge x) + bd(y \wedge y)$$

by multilinearity. Since (3.1) is alternating, $x \wedge x$ and $y \wedge y$ vanish. By (3.2), $y \wedge x = -x \wedge y$. Feeding this into (3.3) gives

$$(ax + by) \wedge (cx + dy) = ad(x \wedge y) - bc(x \wedge y) = (ad - bc)(x \wedge y),$$

so $\Lambda^2(M)$ is spanned by $x \wedge y$ when M is spanned by x and y . That was faster than Example 3.2!

Example 3.7. Suppose M is spanned by three elements e_1, e_2 , and e_3 . We will show $e_1 \wedge e_2$, $e_1 \wedge e_3$, and $e_2 \wedge e_3$ span $\Lambda^2(M)$. We know by the definition of $\Lambda^2(M)$ that $\Lambda^2(M)$ is spanned by all $m \wedge m'$, so it suffices to show every $m \wedge m'$ is a linear combination of $e_1 \wedge e_2$, $e_1 \wedge e_3$, and $e_2 \wedge e_3$. Writing

$$m = ae_1 + be_2 + ce_3, \quad m' = a'e_1 + b'e_2 + c'e_3,$$

the multilinearity and the alternating property ($e_i \wedge e_i = 0$ and $e_i \wedge e_j = -e_j \wedge e_i$ for $i \neq j$) imply

$$\begin{aligned} m \wedge m' &= (ae_1 + be_2 + ce_3) \wedge (a'e_1 + b'e_2 + c'e_3) \\ &= ae_1 \wedge (b'e_2 + c'e_3) + be_2 \wedge (a'e_1 + c'e_3) + ce_3 \wedge (a'e_1 + b'e_2) \\ &= ab'(e_1 \wedge e_2) + ac'(e_1 \wedge e_3) + ba'(e_2 \wedge e_1) + bc'(e_2 \wedge e_3) + \\ &\quad ca'(e_3 \wedge e_1) + cb'(e_3 \wedge e_2) \\ &= (ab' - ba')(e_1 \wedge e_2) + (ac' - ca')(e_1 \wedge e_3) + (bc' - cb')(e_2 \wedge e_3). \end{aligned}$$

If we write this formula with the first and third terms exchanged as

$$(3.4) \quad m \wedge m' = (bc' - cb')(e_2 \wedge e_3) + (ac' - ca')(e_1 \wedge e_3) + (ab' - ba')(e_1 \wedge e_2),$$

it looks quite close to the cross product on \mathbf{R}^3 :

$$(a, b, c) \times (a', b', c') = (bc' - cb', -(ac' - ca'), ab' - ba').$$

(There is a way of making $e_3 \wedge e_1$ the more natural wedge product than $e_1 \wedge e_3$ in $\Lambda^2(\mathbf{R}^3)$, so (3.4) would then match the coordinates of the cross-product everywhere. This uses the Hodge-star operator. We don't discuss that here.)

In a tensor power of a module, every tensor is a sum of elementary tensors but most elements are not themselves elementary tensors. The same thing happens with exterior powers: a general element of $\Lambda^k(M)$ is a sum of elementary wedge products, but is *not* itself of this form.

Example 3.8. Let M be spanned by e_1, e_2, e_3 , and e_4 . The exterior square $\Lambda^2(M)$ is spanned by the pairs

$$e_1 \wedge e_2, e_1 \wedge e_3, e_1 \wedge e_4, e_2 \wedge e_3, e_2 \wedge e_4, e_3 \wedge e_4.$$

When M is free and $\{e_1, e_2, e_3, e_4\}$ is a basis of M , the sum $e_1 \wedge e_2 + e_3 \wedge e_4$ in $\Lambda^2(M)$ is *not* an elementary wedge product: it can't be expressed in the form $m \wedge m'$. We'll see why (in most cases) in Example 8.9. On the other hand, the sum

$$e_1 \wedge e_2 + 3(e_1 \wedge e_3) + 3(e_1 \wedge e_4) + 2(e_2 \wedge e_3) + 2(e_2 \wedge e_4)$$

in $\Lambda^2(M)$ doesn't look like an elementary wedge product but it is! It equals

$$(e_1 + e_2 + e_3 + e_4) \wedge (e_2 + 3e_3 + 3e_4).$$

Check equality by expanding this out using multilinearity and the relations $e_i \wedge e_i = 0$ and $e_i \wedge e_j = -e_j \wedge e_i$ for $i \neq j$.

A linear map out of $\Lambda^k(M)$ is completely determined by its values on the elementary wedge products $m_1 \wedge \cdots \wedge m_k$, since they span the module. But elementary wedge products, like elementary tensors, are not linearly independent, so verifying there is a linear map out of $\Lambda^k(M)$ with some prescribed behavior on all the elementary wedge products has to be done carefully. Proceed by first introducing a function on M^k which is multilinear and alternating whose value at (m_1, \dots, m_k) is what you want the value to be at $m_1 \wedge \cdots \wedge m_k$, and then it automatically factors through $\Lambda^k(M)$ as a linear map with the desired value at $m_1 \wedge \cdots \wedge m_k$. This is like creating homomorphisms out of a quotient group G/N by first making a homomorphism out of G with the desired values and then checking N is in the kernel.

Example 3.9. There is a unique linear map $\Lambda^2(M) \rightarrow M^{\otimes 2}$ such that

$$m_1 \wedge m_2 \mapsto m_1 \otimes m_2 - m_2 \otimes m_1.$$

To construct such a map, start by letting $f: M^2 \rightarrow M^{\otimes 2}$ by $f(m_1, m_2) = m_1 \otimes m_2 - m_2 \otimes m_1$. This is bilinear and $f(m, m) = 0$, so f is alternating and thus there is a linear map $\Lambda^2(M) \rightarrow M^{\otimes 2}$ sending any elementary wedge product $m_1 \wedge m_2$ to $f(m_1, m_2) = m_1 \otimes m_2 - m_2 \otimes m_1$. Since a linear map out of $\Lambda^2(M)$ is determined by its values on elementary wedge products, f is the only linear map with the given values on all $m_1 \wedge m_2$.

Here are some basic questions about exterior powers.

Questions

- (1) What does it mean to say $m_1 \wedge \cdots \wedge m_k = 0$?
- (2) What does it mean to say $\Lambda^k(M) = 0$?

(3) What does it mean to say $m_1 \wedge \cdots \wedge m_k = m'_1 \wedge \cdots \wedge m'_k$?

Answers

(1) Saying $m_1 \wedge \cdots \wedge m_k$ equals 0 means every alternating multilinear map out of M^k vanishes at (m_1, \dots, m_k) . Indeed, since every alternating multilinear map out of M^k induces a linear map out of $\Lambda^k(M)$ which sends $m_1 \wedge \cdots \wedge m_k$ to the same place as (m_1, \dots, m_k) , if $m_1 \wedge \cdots \wedge m_k = 0$ then the linear map out of $\Lambda^k(M)$ must send $m_1 \wedge \cdots \wedge m_k$ to 0 (linear maps send 0 to 0) so the original alternating multilinear map we started with out of M^k has to equal 0 at (m_1, \dots, m_k) . Conversely, if every alternating multilinear map out of M^k sends (m_1, \dots, m_k) to 0 then $m_1 \wedge \cdots \wedge m_k = 0$ because the canonical map $M^k \xrightarrow{\wedge} \Lambda^k(M)$ is a particular example of an alternating multilinear map out of M^k and it sends (m_1, \dots, m_k) to $m_1 \wedge \cdots \wedge m_k$.

Thus you can prove a particular elementary wedge product $m_1 \wedge \cdots \wedge m_k$ is *not* 0 by finding an alternating multilinear map on M^k which is not equal to 0 at (m_1, \dots, m_k) .

(2) To say $\Lambda^k(M) = 0$ means every alternating multilinear map on M^k is identically 0.

To show $\Lambda^k(M) \neq 0$, find an example of an alternating multilinear map on M^k which is not identically 0.

(3) The condition $m_1 \wedge \cdots \wedge m_k = m'_1 \wedge \cdots \wedge m'_k$ means every alternating multilinear map on M^k takes the same values at (m_1, \dots, m_k) and at (m'_1, \dots, m'_k) .

Remark 3.10. Unlike the tensor product, which can be defined between different R -modules, there is no “exterior product” of two unrelated R -modules. This is because the concept of exterior power is bound up with the idea of alternating multilinear functions, and permuting variables in a multivariable function only makes sense when the function has its variables coming from the *same* module.

4. SPANNING SETS FOR EXTERIOR POWERS

Let’s look more closely at spanning sets of an exterior power module. If M is finitely generated (not necessarily free!), with spanning set x_1, \dots, x_d , then any tensor power $M^{\otimes k}$ is finitely generated as an R -module by the d^k tensors $x_{i_1} \otimes \cdots \otimes x_{i_k}$ where $1 \leq i_1, \dots, i_k \leq d$: at first we know $M^{\otimes k}$ is spanned by all the elementary tensors $m_1 \otimes \cdots \otimes m_k$, but write each m_i as an R -linear combination of x_1, \dots, x_d and then expand out using the multilinearity of \otimes to express every elementary tensor in $M^{\otimes k}$ as an R -linear combination of the tensors $x_{i_1} \otimes \cdots \otimes x_{i_k}$. (As a general rule this spanning set for $M^{\otimes k}$ can’t be reduced further: when M is free and x_1, \dots, x_d is a basis then the d^k elementary tensors $x_{i_1} \otimes \cdots \otimes x_{i_k}$ are a basis of $M^{\otimes k}$.) Since $\Lambda^k(M)$ is a quotient module of $M^{\otimes k}$, it is spanned as an R -module by the d^k elementary wedge products $x_{i_1} \wedge \cdots \wedge x_{i_k}$ where $1 \leq i_1, \dots, i_k \leq d$. Thus exterior powers of a finitely generated R -module are finitely generated.

Theorem 4.1. *If M has a d -element spanning set then $\Lambda^k(M) = \{0\}$ for $k > d$.*

For example, $\Lambda^2(R) = 0$, and more generally $\Lambda^k(R^d) = 0$ for $k > d$.

Proof. Let x_1, \dots, x_d span M . When $k > d$, any $x_{i_1} \wedge \cdots \wedge x_{i_k}$ contains two equal terms, so it is zero. Thus $\Lambda^k(M)$ is spanned by 0, so it is 0. \square

There is a lot of redundancy in the d^k elementary wedge products $x_{i_1} \wedge \cdots \wedge x_{i_k}$ coming from a spanning set $\{x_1, \dots, x_d\}$ of M . For instance, by the alternating property such an elementary wedge product vanishes if two terms in it are equal. Therefore we can discard

from our spanning set for $\Lambda^k(M)$ those $x_{i_1} \wedge \cdots \wedge x_{i_k}$ where an x_i appears twice and we are still left with a spanning set. Moreover, by (3.2) two elementary wedge products containing the same factors in different order are equal up to sign, so our spanning set for $\Lambda^k(M)$ as an R -module can be reduced further to the elements $x_{i_1} \wedge \cdots \wedge x_{i_k}$ where the indices are strictly increasing: $1 \leq i_1 < \cdots < i_k \leq d$. The number of such k -tuples of indices is $\binom{d}{k}$. So $\Lambda^k(M)$ has a spanning set of size $\binom{d}{k}$, which may or may not be reducible further. We will now prove that there can be no further reduction when M is free and x_1, \dots, x_d is a basis of M .

Theorem 4.2. *If M is free then $\Lambda^k(M)$ is free, provided $k \leq \text{rank}(M)$ if M has a finite rank and with no constraint on k if M has infinite rank. Explicitly, if $M \neq 0$ is finite free with basis e_1, \dots, e_d , then for $1 \leq k \leq d$ the $\binom{d}{k}$ elementary wedge products*

$$e_{i_1} \wedge \cdots \wedge e_{i_k} \text{ where } 1 \leq i_1 < \cdots < i_k \leq d$$

are a basis of $\Lambda^k(M)$. In particular, $\Lambda^k(M)$ is free of rank $\binom{d}{k}$ for $0 \leq k \leq d$ and $\Lambda^k(M) = 0$ for $k > d$. If M has an infinite basis $\{e_i\}_{i \in I}$ and we put a well-ordering on the index set I , then for any $k \geq 1$ a basis of $\Lambda^k(M)$ is $\{e_{i_1} \wedge \cdots \wedge e_{i_k}\}_{i_1 < i_2 < \cdots < i_k}$.

Theorem 4.2 is the first nontrivial result about exterior powers, as it tells us a situation where exterior powers are guaranteed to be nonzero, and in fact be “as big as possible.” Read the proof closely, as otherwise you may feel somewhat uneasy about exactly why exterior powers of free modules are free.

Proof. For $k = 0$ there is nothing to show. Take $k \geq 1$. The idea in the proof is to embed $\Lambda^k(M)$ as a submodule of $M^{\otimes k}$ and exploit what we know already about $M^{\otimes k}$ for free M . The embedding we will write down may look like it comes out of nowhere, but it is very common in differential geometry and we make some remarks about this after the proof. The basic idea is to turn elementary wedge products into skew-symmetric tensors by an averaging process.

The function $M^k \rightarrow M^{\otimes k}$ given by

$$(m_1, \dots, m_k) \mapsto \sum_{\sigma \in S_k} (\text{sign } \sigma) m_{\sigma(1)} \otimes \cdots \otimes m_{\sigma(k)}$$

is multilinear since each summand contains each m_i once. This function is also alternating. To prove this, by Corollary 2.9 it suffices to check the function vanishes at k -tuples with adjacent equal coordinates. If $m_i = m_{i+1}$ then for each $\sigma \in S_k$ the terms in the sum at σ and $\sigma(i \ i + 1)$ are negatives of each other. Now the universal mapping property of the exterior power tells us there is an R -linear map $\alpha_{k,M}: \Lambda^k(M) \rightarrow M^{\otimes k}$ such that

$$(4.1) \quad \alpha_{k,M}(m_1 \wedge \cdots \wedge m_k) = \sum_{\sigma \in S_k} (\text{sign } \sigma) m_{\sigma(1)} \otimes \cdots \otimes m_{\sigma(k)}.$$

The case $k = 2$, by the way, is Example 3.9.

Although $\alpha_{k,M}$ exists for all M , injectivity of $\alpha_{k,M}$ is not a general property. Our proof of injectivity for free M will use a basis. We’ll write the proof with a finite basis, and the reader can make the minor changes to see the same argument works if M has an infinite basis.

Let M have basis e_1, \dots, e_d . Injectivity of $\alpha_{k,M}: \Lambda^k(M) \rightarrow M^{\otimes k}$ is clear for $k > d$, so we may take $1 \leq k \leq d$. We may even take $k \geq 2$, as the theorem is obvious for $k = 1$. Since

the e_i 's span M as an R -module, the elementary wedge products

$$(4.2) \quad e_{i_1} \wedge \cdots \wedge e_{i_k} \text{ where } 1 \leq i_1 < \cdots < i_k \leq d$$

span $\Lambda^k(M)$. (Here we need that the indices are totally ordered.) We know already that $M^{\otimes k}$ has a *basis*

$$e_{i_1} \otimes \cdots \otimes e_{i_k} \text{ where } 1 \leq i_1, \dots, i_k \leq d,$$

where no inequalities are imposed on the indices.

Suppose $\omega \in \Lambda^k(M)$ satisfies $\alpha_{k,M}(\omega) = 0$. Write

$$\omega = \sum_{1 \leq i_1 < \cdots < i_k \leq d} c_{i_1, \dots, i_k} e_{i_1} \wedge \cdots \wedge e_{i_k}.$$

with $c_{i_1, \dots, i_k} \in R$. Then the condition $\alpha_{k,M}(\omega) = 0$ becomes

$$\sum_{1 \leq i_1 < \cdots < i_k \leq d} c_{i_1, \dots, i_k} \sum_{\sigma \in S_k} (\text{sign } \sigma) e_{i_{\sigma(1)}} \otimes \cdots \otimes e_{i_{\sigma(k)}} = 0,$$

which is the same as

$$\sum_{\sigma \in S_k} \sum_I (\text{sign } \sigma) c_I e_{\sigma(I)} = 0,$$

where I runs over all strictly increasing k -tuples (i_1, \dots, i_k) from 1 to d , with c_I and $e_{\sigma(I)}$ having an obvious meaning in this context. The vectors $\{e_{\sigma(I)}\}_{\sigma, I}$ are a *basis* of $M^{\otimes k}$, so all c_I are 0. This proves $\alpha_{k,M}$ is injective and it also shows (4.2) is a linearly independent subset of $\Lambda^k(M)$, so it is a basis (spans and is linearly independent). \square

Exterior powers are closely connected to determinants, and most proofs of Theorem 4.2 for finite free M use the determinant. What we used in lieu of theorems about determinants is our knowledge of bases of tensor powers of a free module. For aesthetic reasons, we want to come back later and prove properties of the determinant using exterior powers, so we did not use the determinant directly in the proof of Theorem 4.2. However, the linear map $\alpha_{k,M}$ looks a lot like a determinant.

When V is a finite-dimensional vector space, it is free so $\alpha_{k,V}: \Lambda^k(V) \hookrightarrow V^{\otimes k}$ given by (4.1) is an embedding. It means we can think of $\Lambda^k(V)$ as a *subspace* of the tensor power $V^{\otimes k}$ instead of as a quotient space. This viewpoint is widely used in differential geometry, where vector spaces are defined over \mathbf{R} or \mathbf{C} and the image of $\Lambda^k(V)$ in $V^{\otimes k}$ is the subspace of skew-symmetric tensors. The embedding $\alpha_{k,V}$ has an unfortunate scaling problem: when we embed $\Lambda^k(V)$ into $V^{\otimes k}$ with $\alpha_{k,V}$ and then reduce $V^{\otimes k}$ back to $\Lambda^k(V)$ with the canonical map \wedge , the composite map $\Lambda^k(V) \xrightarrow{\alpha_{k,V}} V^{\otimes k} \xrightarrow{\wedge} \Lambda^k(V)$ is *not* the identity map on $\Lambda^k(V)$, but is multiplication by $k!$. We can verify this by checking it on elementary wedge products:

$$\begin{aligned} v_1 \wedge \cdots \wedge v_k &\mapsto \sum_{\sigma \in S_k} (\text{sign } \sigma) v_{\sigma(1)} \otimes \cdots \otimes v_{\sigma(k)} \\ &\mapsto \sum_{\sigma \in S_k} (\text{sign } \sigma) v_{\sigma(1)} \wedge \cdots \wedge v_{\sigma(k)} \\ &= \sum_{\sigma \in S_k} (\text{sign } \sigma) (\text{sign } \sigma) v_1 \wedge \cdots \wedge v_k \\ &= k! v_1 \wedge \cdots \wedge v_k. \end{aligned}$$

This suggests a better embedding of $\Lambda^k(V)$ into $V^{\otimes k}$ is $\frac{1}{k!}\alpha_{k,V}$, which is given by the formula

$$v_1 \wedge \cdots \wedge v_k \mapsto \frac{1}{k!} \sum_{\sigma \in S_k} (\text{sign } \sigma) v_{\sigma(1)} \otimes \cdots \otimes v_{\sigma(k)}.$$

The composite map $\Lambda^k(V) \xrightarrow{(1/k!)\alpha_{k,V}} V^{\otimes k} \xrightarrow{\wedge} \Lambda^k(V)$ is the identity, but this rescaled embedding only makes sense if $k! \neq 0$ in the scalar field. That is fine for real and complex vector spaces (as in differential geometry), but it is not a universal method. So either you can take your embedding $\Lambda^k(V) \hookrightarrow V^{\otimes k}$ using $\alpha_{k,V}$ for all vector spaces and make $\Lambda^k(V) \rightarrow V^{\otimes k} \rightarrow \Lambda^k(V)$ be multiplication by $k!$, or you can have an embedding $\Lambda^k(V) \hookrightarrow V^{\otimes k}$ that only makes sense when $k! \neq 0$. Either way, this mismatch between $\Lambda^k(V)$ as a quotient space of $V^{\otimes k}$ (correct definition) and as a subspace of $V^{\otimes k}$ (incorrect but widely used definition) leads to a lot of excess factorials in formulas when exterior powers are regarded as subspaces of tensor powers instead of as quotient spaces of them.

There are other approaches to the proof of Theorem 4.2 when M is finite free. In [1, pp. 90–91] an explicit finite free R -module is constructed which has the same universal mapping property as $\Lambda^k(M)$, so $\Lambda^k(M)$ has to be finite free by Corollary 3.5. The other aspects of Theorem 4.2 (the rank of $\Lambda^k(M)$ and an explicit basis) can be read off from the proof in [1]. In [4, pp. 747–751], Theorem 4.2 for finite free M is proved using what is called there the Grassmann algebra of M (which we will meet later under the label exterior algebra of M).

When M is free of rank d and $k \leq d$, we will call the basis $\{e_{i_1} \wedge \cdots \wedge e_{i_k} : 1 \leq i_1 < \cdots < i_k \leq d\}$ of $\Lambda^k(M)$ the *corresponding basis* from the choice of basis e_1, \dots, e_d of M .

Example 4.3. If M is free of rank d with basis e_1, \dots, e_d then $\Lambda^2(M)$ is free of rank $\binom{d}{2}$ with corresponding basis $\{e_i \wedge e_j : 1 \leq i < j \leq d\}$.

Remark 4.4. When M is a free R -module, its tensor powers $M^{\otimes k}$ are free R -modules. While $\Lambda^k(M)$ is a quotient module of $M^{\otimes k}$, it does not follow from this alone that $\Lambda^k(M)$ must be free: the quotient of a free module is not generally free (consider R/I where I is a proper nonzero ideal). Work was really needed to show exterior powers of free modules are free modules.

When M is free of rank d , $\Lambda^k(M) \neq 0$ when $k \leq d$ and $\Lambda^k(M) = 0$ when $k > d$. This is why we call $\Lambda^d(M)$ the *top exterior power*. It is free of rank 1; a basis of $\Lambda^d(M)$ is $e_1 \wedge \cdots \wedge e_d$ if e_1, \dots, e_d is a basis of M . Although $\Lambda^d(M)$ is isomorphic to R as an R -module, it is not naturally isomorphic: there is no canonical isomorphism between them.

When M has a d -element spanning set with d minimally chosen, and M is not free, it might happen that $\Lambda^d(M) = 0$. For example, consider a non-principal ideal $I := Rx + Ry$ in R with two generators. The module $\Lambda^2(I)$ is spanned as an R -module by $x \wedge y$.¹ Sometimes $\Lambda^2(I)$ is zero and sometimes it is nonzero.

Example 4.5. Let $R = \mathbf{Z}[\sqrt{-5}]$ and $I = (2, 1 + \sqrt{-5})$. Set $\omega := 2 \wedge (1 + \sqrt{-5}) \in \Lambda^2(I)$. We will show $2\omega = 0$ and $3\omega = 0$, so $\omega = 0$ (just subtract) and thus $\Lambda^2(I) = 0$:

$$\begin{aligned} 2\omega &= 2 \wedge 2(1 + \sqrt{-5}) = (1 + \sqrt{-5})(2 \wedge 2) = 0, \\ 3\omega &= 6 \wedge (1 + \sqrt{-5}) = (1 - \sqrt{-5})((1 + \sqrt{-5}) \wedge (1 + \sqrt{-5})) = 0. \end{aligned}$$

¹It is important to realize $x \wedge y$ here means an elementary wedge product in $\Lambda^2(I)$, *not* in $\Lambda^2(R)$; the latter exterior square is 0 all the time.

Example 4.6. Let $R = A[X, Y]$ be the polynomial ring in two variables over a nonzero commutative ring A . Let $I = (X, Y)$ in R . We will show $X \wedge Y$ in $\Lambda^2(I)$ is nonzero by writing down a linear map out of $\Lambda^2(I)$ whose value on $X \wedge Y$ is nonzero.

Define $B: I^2 \rightarrow A$ to be the determinant on degree one coefficients:

$$B(aX + bY + \cdots, cX + dY + \cdots) = ad - bc.$$

Regard the target module A as an R -module through scaling by the constant term of a polynomial: $f(X, Y)a = f(0, 0)a$. (That is, we basically treat A as $R/I = A[X, Y]/(X, Y)$.) Then B is R -bilinear, and it is alternating too. Since $B(X, Y) = 1$, B induces a linear map $L: \Lambda^2(I) \rightarrow A$ where $L(X \wedge Y) = B(X, Y) = 1$, so $\Lambda^2(I) \neq 0$.

Remark 4.7. An analogue of Example 4.5 in the real quadratic ring $\mathbf{Z}[\sqrt{5}]$ has a different result. For $J = (2, 1 + \sqrt{5})$, $2 \wedge (1 + \sqrt{5})$ in $\Lambda^2(J)$ is not 0, so $\Lambda^2(J) \neq 0$. Constructing an alternating bilinear map out of $J \times J$ that is not identically 0 can be done using the ideas in Example 4.6, and details are left to the reader (Hint: Find a basis for J as a \mathbf{Z} -module.)

The moral from these two examples is that for nonfree M , the highest k for which $\Lambda^k(M) \neq 0$ need not be the size of a minimal spanning set for the module. It only gives an upper bound, when M is finitely generated.²

An important distinction to remember between tensor and exterior powers is that exterior powers are not recursively defined. Whereas $M^{\otimes(k+1)} \cong M \otimes_R M^{\otimes k}$, we can't say that $\Lambda^{k+1}(M)$ is a product of M and $\Lambda^k(M)$; later on (Section 9) we will introduce the exterior algebra, in which something like this does make sense. To appreciate the lack of a recursive definition of exterior powers, consider the following problem. When M has a d -element spanning set, $\Lambda^k(M) = 0$ for $k \geq d + 1$ by Theorem 4.1. Treating M as $\Lambda^1(M)$, we pose a generalization: if $\Lambda^i(M)$ for some $i > 1$ has a d -element spanning set, is $\Lambda^k(M) = 0$ for $k \geq d + i$? The next theorem settles the $d = 0$ case in the affirmative.

Theorem 4.8. *If $\Lambda^i(M) = 0$ for some $i \geq 1$ then $\Lambda^j(M) = 0$ for all $j \geq i$.*

Proof. It suffices to show $\Lambda^i(M) = 0 \Rightarrow \Lambda^{i+1}(M) = 0$, as then we are done by induction. To prove $\Lambda^{i+1}(M) = 0$ we show all $(i + 1)$ -fold elementary wedge products

$$(4.3) \quad m_1 \wedge \cdots \wedge m_i \wedge m_{i+1}$$

are 0.

First we give a fake proof. Since $\Lambda^i(M) = 0$, $m_1 \wedge \cdots \wedge m_i = 0$, so (4.3) equals $0 \wedge m_{i+1} = 0$. What makes this absurd, at our present level of understanding, is that there is no sense (yet) in which the notation \wedge is “associative,” as the notation \wedge is really just a placeholder to tell us where things go. We can't treat the piece $m_1 \wedge \cdots \wedge m_i$ in (4.3) as its own elementary wedge product having any kind of relation to (4.3). This is like calculus, where students are warned that the separate parts of dy/dx do not have an independent meaning, although later they may learn otherwise, as we too will learn otherwise about \wedge in Section 9.

Now we give a real proof, which in fact contains the germ of the idea in Section 9 to make \wedge into a genuine operation and not just a placeholder. For each elementary wedge product (4.3), which belongs to $\Lambda^{i+1}(M)$, we will create a linear map $\Lambda^i(M) \rightarrow \Lambda^{i+1}(M)$ with (4.3) in its image. Then since $\Lambda^i(M) = 0$, so a linear map out of $\Lambda^i(M)$ has image 0, (4.3) is 0.

²Since $(\mathbf{Q}/\mathbf{Z}) \otimes_{\mathbf{Z}} (\mathbf{Q}/\mathbf{Z}) = 0$, $\Lambda^2(\mathbf{Q}/\mathbf{Z}) = 0$ where we regard \mathbf{Q}/\mathbf{Z} as a \mathbf{Z} -module, so the vanishing of $\Lambda^k(M)$ for some k does not force M to be finitely generated.

Consider the function $M^i \rightarrow \Lambda^{i+1}(M)$ given by

$$(x_1, \dots, x_i) \mapsto x_1 \wedge \cdots \wedge x_i \wedge m_{i+1}.$$

This is multilinear and alternating, so by the universal mapping property of exterior powers there is a linear map $\Lambda^i(M) \rightarrow \Lambda^{i+1}(M)$ where

$$x_1 \wedge \cdots \wedge x_i \mapsto x_1 \wedge \cdots \wedge x_i \wedge m_{i+1}.$$

The left side is 0 for all choices of x_1, \dots, x_i in M , so the right side is 0 for all such choices too. In particular, (4.3) is 0.

Here is a different proof. To say $\Lambda^i(M) = 0$ means any alternating multilinear function out of M^i is identically 0. If $\varphi: M^{i+1} \rightarrow N$ is an alternating multilinear function, and $(m_1, \dots, m_i, m_{i+1}) \in M^{i+1}$, consider the function $\varphi(x_1, \dots, x_i, m_{i+1})$ in x_1, \dots, x_i . It is alternating multilinear in i variables from M . Therefore $\varphi(x_1, \dots, x_i, m_{i+1}) = 0$ for all x_1, \dots, x_i in M , so $\varphi(m_1, \dots, m_i, m_{i+1}) = 0$. \square

Returning to the general question, where $\Lambda^i(M)$ has a d -element spanning set, asking if $\Lambda^k(M) = 0$ for $k \geq d+i$ is the same as asking if $\Lambda^{d+i}(M) = 0$ by Theorem 4.8. The answer is “yes” although more technique is needed for that than we will develop here.

5. EXTERIOR POWERS OF LINEAR MAPS

Having constructed exterior powers of modules, we extend the construction to linear maps between modules. First recall any linear map $\varphi: M \rightarrow N$ between two R -modules induces a linear map $\varphi^{\otimes k}: M^{\otimes k} \rightarrow N^{\otimes k}$ on the k th tensor powers, for any positive integer k , which has the effect

$$\varphi^{\otimes k}(m_1 \otimes \cdots \otimes m_k) = \varphi(m_1) \otimes \cdots \otimes \varphi(m_k)$$

on elementary tensors.

Theorem 5.1. *Let $\varphi: M \rightarrow N$ be a linear map of R -modules. Then for each $k \geq 2$ there is a unique linear map $\wedge^k(\varphi): \Lambda^k(M) \rightarrow \Lambda^k(N)$ with the effect*

$$m_1 \wedge \cdots \wedge m_k \mapsto \varphi(m_1) \wedge \cdots \wedge \varphi(m_k)$$

on all elementary wedge products. For a second linear map $\psi: N \rightarrow P$, $\wedge^k(\psi \circ \varphi) = \wedge^k(\psi) \circ \wedge^k(\varphi)$. Moreover, $\wedge^k(\text{id}_M) = \text{id}_{\Lambda^k(M)}$.

Proof. There is at most one such linear map $\Lambda^k(M) \rightarrow \Lambda^k(N)$ since the elementary wedge products span $\Lambda^k(M)$. To show there is such a linear map, start by backing up and defining a function $f: M^k \rightarrow \Lambda^k(N)$ by

$$f(m_1, \dots, m_k) = \varphi(m_1) \wedge \cdots \wedge \varphi(m_k).$$

This is a multilinear map which is alternating, so by the universal mapping property of the k th exterior power there is a linear map $\Lambda^k(M) \rightarrow \Lambda^k(N)$ with the effect

$$m_1 \wedge \cdots \wedge m_k \mapsto f(m_1, \dots, m_k) = \varphi(m_1) \wedge \cdots \wedge \varphi(m_k).$$

$$\begin{array}{ccc} M^k & \xrightarrow{f} & \Lambda^k(N) \\ \wedge \downarrow & \nearrow & \\ \Lambda^k(M) & & \end{array}$$

This proves the existence of the linear map $\wedge^k(\varphi)$ we are seeking.

To show $\wedge^k(\psi \circ \varphi) = \wedge^k(\psi) \circ \wedge^k(\varphi)$, it suffices since both sides are linear to check that both sides have the same value on each elementary wedge product in $\Lambda^k(M)$. At any elementary wedge product $m_1 \wedge \cdots \wedge m_k$, the left side and right side have the common value $\psi(\varphi(m_1)) \wedge \cdots \wedge \psi(\varphi(m_k))$. That $\wedge^k(\text{id}_M) = \text{id}_{\Lambda^k(M)}$ is easy: $\wedge^k(\text{id}_M)$ is linear and fixes every $m_1 \wedge \cdots \wedge m_k$ and these span $\Lambda^k(M)$, so $\wedge^k(\text{id}_M)$ fixes everything. \square

Theorem 5.1 is also true for $k = 0$ and $k = 1$ by setting $\wedge^0(\varphi) = \text{id}_R$ and $\wedge^1(\varphi) = \varphi$. Recall $\Lambda^0(M) = R$ and $\Lambda^1(M) = M$.

That the passage from φ to $\wedge^k(\varphi)$ respects composition and sends the identity map on a module to the identity map on its k th exterior power is called *functoriality* of the k th exterior power.

Armed with bases for exterior powers of finite free modules, we can write down matrices for exterior powers of linear maps between them. When M and N are finite free of respective ranks m and n , a choice of bases of M and N turns any linear map $\varphi: M \rightarrow N$ into an $n \times m$ matrix. Using the corresponding bases on $\Lambda^k(M)$ and $\Lambda^k(N)$, we can write $\wedge^k(\varphi)$ as an $\binom{n}{k} \times \binom{m}{k}$ matrix. If we want to look at an example, we need to keep m and n small or we will face very large matrices.

Example 5.2. If M and N are free of rank 5 and $\varphi: M \rightarrow N$ is linear, then $\wedge^2(\varphi)$ is represented by a 10×10 matrix since $\binom{5}{2} = 10$.

Example 5.3. Let $L: \mathbf{R}^3 \rightarrow \mathbf{R}^3$ be the linear map given by the matrix

$$(5.1) \quad \begin{pmatrix} 0 & 2 & 0 \\ 1 & 1 & 1 \\ 0 & 3 & 2 \end{pmatrix}.$$

We will compute the matrix for $\wedge^2(L): \Lambda^2(\mathbf{R}^3) \rightarrow \Lambda^2(\mathbf{R}^3)$ with respect to the basis $e_1 \wedge e_2, e_1 \wedge e_3, e_2 \wedge e_3$, where the e_i 's are the standard basis of \mathbf{R}^3 . Going in order,

$$\begin{aligned} \wedge^2(L)(e_1 \wedge e_2) &= L(e_1) \wedge L(e_2) \\ &= e_2 \wedge (2e_1 + e_2 + 3e_3) \\ &= -2(e_1 \wedge e_2) + 3(e_2 \wedge e_3), \end{aligned}$$

$$\begin{aligned} \wedge^2(L)(e_1 \wedge e_3) &= L(e_1) \wedge L(e_3) \\ &= e_2 \wedge (e_2 + 2e_3) \\ &= 2(e_2 \wedge e_3), \end{aligned}$$

and

$$\begin{aligned} \wedge^2(L)(e_2 \wedge e_3) &= L(e_2) \wedge L(e_3) \\ &= (2e_1 + e_2 + 3e_3) \wedge (e_2 + 2e_3) \\ &= 2(e_1 \wedge e_2) + 4(e_1 \wedge e_3) - e_2 \wedge e_3. \end{aligned}$$

Therefore the matrix for $\wedge^2(L)$ relative to this ordered basis is

$$\begin{pmatrix} -2 & 0 & 2 \\ 0 & 0 & 4 \\ 3 & 2 & -1 \end{pmatrix}.$$

Theorem 5.4. Let $\varphi: M \rightarrow N$ be linear. If φ is an R -module isomorphism then $\wedge^k(\varphi)$ is an R -module isomorphism for every k . If φ is surjective then every $\wedge^k(\varphi)$ is surjective.

Proof. It is clear for $k = 0$ and $k = 1$. Let $k \geq 2$. Suppose φ is an isomorphism of R -modules, with inverse $\psi: N \rightarrow M$. Then $\varphi \circ \psi = \text{id}_N$ and $\psi \circ \varphi = \text{id}_M$, so by Theorem 5.1 we have $\wedge^k(\varphi) \circ \wedge^k(\psi) = \wedge^k(\text{id}_N) = \text{id}_{\wedge^k(N)}$ and similarly $\wedge^k(\psi) \circ \wedge^k(\varphi) = \text{id}_{\wedge^k(M)}$.

If φ is surjective, then $\wedge^k(\varphi)$ is surjective because $\wedge^k(N)$ is spanned by the elementary wedge products $n_1 \wedge \cdots \wedge n_k$ for all $n_i \in N$ and these are in the image of $\wedge^k(\varphi)$ explicitly: writing $n_i = \varphi(m_i)$, the elementary wedge product of the n_i 's in $\wedge^k(N)$ is $\varphi(m_1 \wedge \cdots \wedge m_k)$. Since the image of the linear map $\wedge^k(\varphi)$ is a submodule of $\wedge^k(N)$ which contains a spanning set for $\wedge^k(N)$, the image is all of $\wedge^k(N)$. \square

As with tensor products of linear maps, it is *false* that $\wedge^k(\varphi)$ has to be injective if φ is injective.

Example 5.5. Let $R = A[X, Y]$ with A a nonzero commutative ring and let $I = (X, Y)$. In our discussion of tensor products, it was seen that the inclusion map $i: I \rightarrow R$ is injective while its induced R -linear map $i^{\otimes 2}: I^{\otimes 2} \rightarrow R^{\otimes 2} \cong R$ is not injective. Therefore it should come as no surprise that the map $\wedge^2(i): \wedge^2(I) \rightarrow \wedge^2(R)$ also is not injective. Indeed, $\wedge^2(R) = 0$ and we saw in Example 4.6 that $\wedge^2(I) \neq 0$ because there is an R -linear map $L: \wedge^2(I) \rightarrow A$ where $L(X \wedge Y) = 1$.

(In $\wedge^2(I)$ we have $X \wedge Y \neq 0$ while in $\wedge^2(R)$ we have $X \wedge Y = XY(1 \wedge 1) = 0$. There is nothing inconsistent about this, even though $I \subset R$, because the natural map $\wedge^2(i): \wedge^2(I) \rightarrow \wedge^2(R)$ is not injective. The $X \wedge Y$'s in $\wedge^2(I)$ and $\wedge^2(R)$ lie in different modules, and although $\wedge^2(i)$ sends $X \wedge Y$ in the first module to $X \wedge Y$ in the second, linear maps can send a nonzero element to 0 and that is what is happening.)

We now show the linear map $L: \wedge^2(I) \rightarrow A$ is actually an isomorphism of R -modules. In $\wedge^2(I) = \{r(X \wedge Y) : r \in R\}$, $X \wedge Y$ is killed by multiplication by X and Y since $X(X \wedge Y) = X \wedge XY = Y(X \wedge X) = 0$ and likewise for $Y(X \wedge Y)$. So $f(X, Y)(X \wedge Y) = f(0, 0)(X \wedge Y)$ in $\wedge^2(I)$, which means every element of $\wedge^2(I)$ has the form $a(X \wedge Y)$ for some $a \in A$. Thus the function $L': A \rightarrow \wedge^2(I)$ given by $L'(a) = a(X \wedge Y)$ is R -linear and is an inverse to L .

The isomorphism $\wedge^2(I) \cong A$ generalizes to the polynomial ring $R = A[X_1, \dots, X_n]$ for any $n \geq 2$: the ideal $I = (X_1, \dots, X_n)$ in R has $\wedge^n(I) \cong A \cong R/I$, so the inclusion $i: I \rightarrow R$ is injective but $\wedge^n(i): \wedge^n(I) \rightarrow \wedge^n(R) = 0$ is not injective.

Although exterior powers don't preserve injectivity of linear maps in general, there are some cases when they do. This is the topic of the rest of this section. A number of ideas here and later are taken from [2].

Theorem 5.6. *Suppose $\varphi: M \rightarrow N$ is injective and the image $\varphi(M) \subset N$ is a direct summand: $N = \varphi(M) \oplus P$ for some submodule P of N . Then $\wedge^k(\varphi)$ is injective for all $k \geq 0$ and $\wedge^k(M)$ is isomorphic to a direct summand of $\wedge^k(N)$.*

Proof. The result is trivial for $k = 0$ and $k = 1$. Suppose $k \geq 2$.

We will use the splitting criteria for short exact sequences of modules. Since $N = \varphi(M) \oplus P$, we have a linear map $\psi: N \rightarrow M$ which undoes the effect of φ : let $\psi(\varphi(m) + p) = m$. Then $\psi(\varphi(m)) = m$ for all $m \in M$, so

$$(5.2) \quad \psi \circ \varphi = \text{id}_M.$$

(The composite in the other direction, $\varphi \circ \psi$, is definitely not id_N unless $P = 0$, but this does not matter.) Applying \wedge^k to (5.2) gives us linear maps $\wedge^k(\varphi): \wedge^k(M) \rightarrow \wedge^k(N)$ and $\wedge^k(\psi): \wedge^k(N) \rightarrow \wedge^k(M)$ with

$$\wedge^k(\psi) \circ \wedge^k(\varphi) = \wedge^k(\psi \circ \varphi) = \wedge^k(\text{id}_M) = \text{id}_{\wedge^k(M)}$$

by functoriality (Theorem 5.1). In particular, if $\wedge^k(\varphi)(\omega) = 0$ for some $\omega \in \Lambda^k(M)$, then applying $\wedge^k(\psi)$ to both sides gives us $\omega = \wedge^k(\psi)(0) = 0$, so $\wedge^k(\varphi)$ has kernel 0 and thus is injective.

For the short exact sequence $0 \rightarrow \Lambda^k(M) \xrightarrow{\wedge^k(\varphi)} \Lambda^k(N) \rightarrow \Lambda^k(N)/\Lambda^k(M) \rightarrow 0$, the fact that $\wedge^k(\psi)$ is a left inverse to $\wedge^k(\varphi)$ implies by the splitting criteria for short exact sequences that $\Lambda^k(N) \cong \Lambda^k(M) \oplus (\Lambda^k(N)/\Lambda^k(M))$, so $\Lambda^k(M)$ is isomorphic to a direct summand of $\Lambda^k(N)$. \square

Example 5.7. For any R -modules M and M' , the inclusion $i: M \rightarrow M \oplus M'$ where $i(m) = (m, 0)$ has image $M \oplus 0$, which is a direct summand of $M \oplus M'$, so the induced linear map $\Lambda^k(M) \rightarrow \Lambda^k(M \oplus M')$ sending $m_1 \wedge \cdots \wedge m_k$ to $(m_1, 0) \wedge \cdots \wedge (m_k, 0)$ is one-to-one.

Remark 5.8. It is instructive to check that the hypothesis of Theorem 5.6 does not apply to the inclusion $i: (X, Y) \rightarrow A[X, Y]$ in Example 5.5 (which must be so because $\wedge^2(i)$ is not injective). We can write $A[X, Y] = A \oplus (X, Y)$, so (X, Y) is a direct summand of $A[X, Y]$ as abelian groups, or even as A -modules, but this is not a direct sum of $A[X, Y]$ -modules: A is not an ideal in $A[X, Y]$.

Corollary 5.9. *Let K be a field and V and W be K -vector spaces. If the K -linear map $\varphi: V \rightarrow W$ is injective then $\wedge^k(\varphi): \Lambda^k(V) \rightarrow \Lambda^k(W)$ is injective for all $k \geq 0$.*

Proof. The subspace $\varphi(V) \subset W$ is a direct summand of W : pick a basis of $\varphi(V)$ over K , extend it to a basis of the whole space W , and let P be the span of the new part of this full basis: $W = \varphi(V) \oplus P$. Thus the hypothesis of Theorem 5.6 applies to this situation. \square

When working with linear maps of vector spaces (not necessarily finite-dimensional), we have shown

$$\begin{aligned} \varphi: V \rightarrow W \text{ injective} &\implies \wedge^k(\varphi) \text{ injective for all } k \text{ (Corollary 5.9),} \\ \varphi: V \rightarrow W \text{ surjective} &\implies \wedge^k(\varphi) \text{ surjective for all } k \text{ (Theorem 5.4),} \\ \varphi: V \rightarrow W \text{ an isomorphism} &\implies \wedge^k(\varphi) \text{ an isomorphism for all } k. \end{aligned}$$

Replacing vector spaces with modules, the second and third properties are true but the first one may fail (Example 5.5 with $R = A[X, Y]$ and $k = 2$). The first property does remain true for free modules, however.

Theorem 5.10. *Suppose M and N are free R -modules. If a linear map $\varphi: M \rightarrow N$ is injective then $\wedge^k(\varphi): \Lambda^k(M) \rightarrow \Lambda^k(N)$ is injective for all k .*

Notice the free hypothesis! We can't use Corollary 5.9 here (if R is not a field), as the image of φ need not be a direct summand of N . That is, a submodule of a free module is often not a direct summand (unless R is a field). We are not assuming M and N have finite bases.

Proof. The diagram

$$\begin{array}{ccc} \Lambda^k(M) & \xrightarrow{\alpha_{k,M}} & M^{\otimes k} \\ \wedge^k(\varphi) \downarrow & & \downarrow \varphi^{\otimes k} \\ \Lambda^k(N) & \xrightarrow{\alpha_{k,N}} & N^{\otimes k} \end{array}$$

commutes, where the top and bottom maps come from Theorem 4.2. The explicit effect in this diagram on an elementary wedge product in $\Lambda^k(M)$ is given by

$$\begin{array}{ccc} m_1 \wedge \cdots \wedge m_k & \xrightarrow{\alpha_{k,M}} & \sum_{\sigma \in S_k} (\text{sign } \sigma) m_{\sigma(1)} \otimes \cdots \otimes m_{\sigma(k)} \\ \downarrow \wedge^k(\varphi) & & \downarrow \varphi^{\otimes k} \\ \varphi(m_1) \wedge \cdots \wedge \varphi(m_k) & \xrightarrow{\alpha_{k,N}} & \sum_{\sigma \in S_k} (\text{sign } \sigma) \varphi(m_{\sigma(1)}) \otimes \cdots \otimes \varphi(m_{\sigma(k)}) \end{array}$$

Since $\alpha_{k,M}$ and $\alpha_{k,N}$ are injective by Theorem 4.2, and $\varphi^{\otimes k}$ is injective from our development of the tensor product (any tensor power of an injective linear map of *flat* modules is injective, and free modules are flat), $\wedge^k(\varphi)$ has to be injective from commutativity of the diagram. \square

Here is a nice application of Theorem 5.10 and the nonvanishing of $\Lambda^k(R^n)$ for $k \leq n$ (but not $k > n$).

Corollary 5.11. *If $\varphi: R^m \rightarrow R^n$ is a linear map, then surjectivity of φ implies $m \geq n$ and injectivity of φ implies $m \leq n$.*

Proof. First suppose φ is onto. Taking n th exterior powers, $\wedge^n(\varphi): \Lambda^n(R^m) \rightarrow \Lambda^n(R^n)$ is onto by Theorem 5.4. Since $\Lambda^n(R^n) \neq 0$, $\Lambda^n(R^m) \neq 0$, so $n \leq m$.

Now suppose φ is one-to-one. Taking m th exterior powers, $\wedge^m(\varphi): \Lambda^m(R^m) \rightarrow \Lambda^m(R^n)$ is one-to-one by Theorem 5.10, so the nonvanishing of $\Lambda^m(R^m)$ implies $\Lambda^m(R^n) \neq 0$, so $m \leq n$. \square

The proof of Corollary 5.11 is short, but if we unravel it we see that the injectivity part of Corollary 5.11 is a deeper result than the surjectivity, because exterior powers of linear maps don't preserve injectivity in general (Example 5.5). There will be another interesting application of Theorem 5.10 in Section 7 (Theorem 7.4).

Corollary 5.12. *If M is a free module and $\{m_1, \dots, m_s\}$ is a finite linearly independent subset then for any $k \leq s$ the $\binom{s}{k}$ elementary wedge products*

$$(5.3) \quad m_{i_1} \wedge \cdots \wedge m_{i_k} \text{ where } 1 \leq i_1 < \cdots < i_k \leq s$$

are linearly independent in $\Lambda^k(M)$.

Proof. We have an embedding $R^s \hookrightarrow M$ by $\sum_{i=1}^s r_i e_i \mapsto \sum_{i=1}^s r_i m_i$. Since R^s and M are free, the k th exterior power of this linear map is an embedding $\Lambda^k(R^s) \hookrightarrow \Lambda^k(M)$ which sends the basis

$$e_{i_1} \wedge \cdots \wedge e_{i_k}$$

of $\Lambda^k(R^s)$, where $1 \leq i_1 < \cdots < i_k \leq s$, to the elementary wedge products in (5.3), so they are linearly independent in $\Lambda^k(M)$. \square

Remark 5.13. In a vector space, any linearly independent subset extends to a basis. The corresponding result in R^n is generally false. In fact \mathbf{Z}^2 already provides counterexamples: the vector $(2, 2)$ is linearly independent by itself but can't belong to a basis because a basis vector in \mathbf{Z}^2 must have relatively prime coordinates. Since any linearly independent subset of R^n has at most n terms in it, by Corollary 5.11, it is natural to ask if every maximal linearly independent subset of R^n has n vectors in it (which need not be a basis). This is true in \mathbf{Z}^2 , *e.g.*, $(2, 2)$ is part of the linearly independent subset $\{(2, 2), (1, 0)\}$.

However, it is *not* true in general that every maximal linearly independent subset of R^n has n vectors in it. There are rings R such that R^2 contains a vector v such that $\{v\}$ is

linearly independent (meaning it's torsion-free) but there is no linearly independent subset $\{v, w\}$ in R^2 . An example, due to David Speyer, is the following: let R be the ring of functions $\mathbf{C}^2 - \{(0, 0)\} \rightarrow \mathbf{C}$ which coincide with a polynomial function at all but finitely many points. (The finitely many exceptional points can vary.) Letting z and w be the coordinate functions on \mathbf{C}^2 , in R^2 the vector (z, w) is linearly independent and is not part of any larger linearly independent subset.

6. DETERMINANTS

Now we put exterior powers to work in the development of the determinant, whose properties have up until now played no role except for a 2×2 determinant in Example 5.5.

Let M be a free R -module of rank $d \geq 1$. The top exterior power $\Lambda^d(M)$ is a free R -module of rank 1, so any linear map $\Lambda^d(M) \rightarrow \Lambda^d(M)$ is scaling by an element of R . For a linear map $\varphi: M \rightarrow M$, the induced linear map $\Lambda^d(\varphi): \Lambda^d(M) \rightarrow \Lambda^d(M)$ is scaling by what element of R ?

Theorem 6.1. *If M is a free R -module of rank $d \geq 1$ and $\varphi: M \rightarrow M$ is a linear map, its top exterior power $\Lambda^d(\varphi): \Lambda^d(M) \rightarrow \Lambda^d(M)$ is multiplication by $\det \varphi \in R$.*

Proof. We want to show $\Lambda^d(\varphi)(\omega) = (\det \varphi)\omega$ for all $\omega \in \Lambda^d(M)$. It suffices to check this when ω is a basis of $\Lambda^d(M)$. Let e_1, \dots, e_d be a basis for M , so $e_1 \wedge \dots \wedge e_d$ is a basis for $\Lambda^d(M)$. We will show

$$\Lambda^d(\varphi)(e_1 \wedge \dots \wedge e_d) = (\det \varphi)(e_1 \wedge \dots \wedge e_d).$$

By definition,

$$\Lambda^d(\varphi)(e_1 \wedge \dots \wedge e_d) = \varphi(e_1) \wedge \dots \wedge \varphi(e_d).$$

Let $\varphi(e_j) = \sum_{i=1}^d a_{ij}e_i$. Then (a_{ij}) is the matrix representation for φ in the ordered basis e_1, \dots, e_d and

$$\begin{aligned} \Lambda^d(\varphi)(e_1 \wedge \dots \wedge e_d) &= \sum_{i=1}^d a_{i1}e_i \wedge \dots \wedge \sum_{i=1}^d a_{id}e_i \\ &= \sum_{i_1=1}^d a_{i_1 1}e_{i_1} \wedge \dots \wedge \sum_{i_d=1}^d a_{i_d d}e_{i_d}, \end{aligned}$$

where we introduce different labels for the summation indices because we are about to combine terms using multilinearity:

$$\Lambda^d(\varphi)(e_1 \wedge \dots \wedge e_d) = \sum_{i_1, \dots, i_d=1}^d a_{i_1 1} \dots a_{i_d d} e_{i_1} \wedge \dots \wedge e_{i_d}.$$

In this sum, terms with equal indices can be dropped (the wedge product vanishes) so all we are left with is a sum over d -tuples of distinct indices. Since d distinct integers from 1 to d must be $1, 2, \dots, d$ in some rearrangement, we can write $i_1 = \sigma(1), \dots, i_d = \sigma(d)$ as σ runs over S_d :

$$\Lambda^d(\varphi)(e_1 \wedge \dots \wedge e_d) = \sum_{\sigma \in S_d} a_{\sigma(1)1} \dots a_{\sigma(d)d} (e_{\sigma(1)} \wedge \dots \wedge e_{\sigma(d)}).$$

By (3.2), this becomes

$$\wedge^d(\varphi)(e_1 \wedge \cdots \wedge e_d) = \sum_{\sigma \in S_d} (\text{sign } \sigma) a_{\sigma(1)1} \cdots a_{\sigma(d)d} (e_1 \wedge \cdots \wedge e_d).$$

Thus $\wedge^d(\varphi)$ acts on the 1-element basis $e_1 \wedge \cdots \wedge e_d$ of $\Lambda^d(M)$ as multiplication by the number we recognize as $\det((a_{ij})^\top) = \det(a_{ij}) = \det(\varphi)$, so it acts on every element of $\Lambda^d(M)$ as scaling by $\det(\varphi)$. \square

Since we did not use determinants before, we could *define* the determinant of a linear operator φ on a (nonzero) finite free module M to be the scalar by which φ acts on the top exterior power of M . Then the proof of Theorem 6.1 shows $\det \varphi$ can be computed from any matrix representation of φ by the usual formula, and it shows this formula is independent of the choice of matrix representation for φ since our construction of exterior powers was coordinate-free. Since $\wedge^k(\text{id}_M) = \text{id}_{\Lambda^k(M)}$, the determinant of the identity map is 1. Here is a slick proof that the determinant is multiplicative:

Corollary 6.2. *If M is a nonzero finite free R -module and φ and ψ are linear maps $M \rightarrow M$, then $\det(\psi \circ \varphi) = \det(\psi) \det(\varphi)$.*

Proof. Let M have rank $d \geq 1$. By Theorem 5.1, $\wedge^d(\psi \circ \varphi) = \wedge^d(\psi) \circ \wedge^d(\varphi)$. Both sides are linear maps $\Lambda^d(M) \rightarrow \Lambda^d(M)$ and $\Lambda^d(M)$ is free of rank 1. The left side is multiplication by $\det(\psi \circ \varphi)$. The right side is multiplication by $\det(\varphi)$ followed by multiplication by $\det(\psi)$, which is multiplication by $\det(\psi) \det(\varphi)$. Thus, by checking both sides on a one-element basis of $\Lambda^d(M)$, we obtain $\det(\psi \circ \varphi) = \det(\psi) \det(\varphi)$. \square

Continuing a purely logical development (not assuming prior knowledge of determinants, that is), at this point we could introduce the characteristic polynomial and prove the Cayley-Hamilton theorem. One of the corollaries of the Cayley-Hamilton theorem is that $\text{GL}_d(R) = \{A \in \text{M}_d(R) : \det A \in R^\times\}$. That is used in the next result, which characterizes bases in R^d using $\Lambda^d(R)$.

Corollary 6.3. *Let M be a free R -module of rank $d \geq 1$. For $x_1, \dots, x_d \in M$, $\{x_1, \dots, x_d\}$ is a basis of M if and only if $x_1 \wedge \cdots \wedge x_d$ is a basis of $\Lambda^d(M)$.*

Proof. We know by Theorem 4.2 that if $\{x_1, \dots, x_d\}$ is a basis of M then $x_1 \wedge \cdots \wedge x_d$ is a basis of $\Lambda^d(M)$. We now want to go the other way: if $x_1 \wedge \cdots \wedge x_d$ is a basis of $\Lambda^d(M)$ we show $\{x_1, \dots, x_d\}$ is a basis of M .

Since M is free of rank d it has some basis, say $\{e_1, \dots, e_d\}$. Write the x_j 's in terms of this basis: $x_j = \sum_{i=1}^d a_{ij} e_i$, where $a_{ij} \in R$. That means the linear map $A: M \rightarrow M$ given by $A(e_j) = x_j$ for all j has matrix representation (a_{ij}) in the basis $\{e_1, \dots, e_d\}$. Therefore

$$x_1 \wedge \cdots \wedge x_d = A e_1 \wedge \cdots \wedge A e_d = \wedge^d(A)(e_1 \wedge \cdots \wedge e_d) = (\det A)(e_1 \wedge \cdots \wedge e_d).$$

Since $x_1 \wedge \cdots \wedge x_d$ and $e_1 \wedge \cdots \wedge e_d$ are both bases of $\Lambda^d(M)$ (the first by hypothesis and the second by Theorem 4.2), the scalar by which they differ must be a unit: $\det A \in R^\times$. Therefore $A \in \text{GL}_d(R)$, so the x_j 's must be a basis of R^d because $A(e_j) = x_j$ and the e_j 's are a basis. \square

While the top exterior power of a linear operator on a finite free module is multiplication by its determinant, the lower-order exterior powers of a linear map between finite free modules have matrix representations whose entries are determinants. Let $\varphi: M \rightarrow N$ be

linear, with M and N finite free of positive ranks m and n , respectively. Take $k \leq m$ and $k \leq n$ since otherwise $\Lambda^k(M)$ or $\Lambda^k(N)$ is 0. Pick bases e_1, \dots, e_m for M and f_1, \dots, f_n for N . The corresponding basis of $\Lambda^k(M)$ is all $e_{j_1} \wedge \dots \wedge e_{j_k}$ where

$$1 \leq j_1 < \dots < j_k \leq m,$$

Similarly, the corresponding basis for $\Lambda^k(N)$ is all $f_{i_1} \wedge \dots \wedge f_{i_k}$ where $1 \leq i_1 < \dots < i_k \leq n$. The matrix for $\Lambda^k(\varphi)$ relative to these bases of $\Lambda^k(M)$ and $\Lambda^k(N)$ is therefore naturally indexed by pairs of increasing k -tuples.

Theorem 6.4. *With notation as above, let $[\varphi]$ denotes the $n \times m$ matrix for φ relative to the choice of bases for M and N . Relative to the corresponding bases on $\Lambda^k(M)$ and $\Lambda^k(N)$, the matrix entry for $\Lambda^k(\varphi)$ in row position (i_1, \dots, i_k) and column position (j_1, \dots, j_k) is the determinant of the $k \times k$ matrix built from rows i_1, \dots, i_k and columns j_1, \dots, j_k of $[\varphi]$.*

Proof. The matrix entry in question is the coefficient of $f_{i_1} \wedge \dots \wedge f_{i_k}$ in the expansion of $\Lambda^k(\varphi)(e_{j_1} \wedge \dots \wedge e_{j_k}) = \varphi(e_{j_1}) \wedge \dots \wedge \varphi(e_{j_k})$. Details are left to the reader. \square

In short, this says the k th exterior power of a linear map φ has matrix entries which are $k \times k$ determinants of submatrices of the matrix of φ .

Example 6.5. In Example 5.3 we computed the second exterior power of the linear map $\mathbf{R}^3 \rightarrow \mathbf{R}^3$ given by the 3×3 matrix L in (5.1). The result is a matrix of size $\binom{3}{2} \times \binom{3}{2} = 3 \times 3$ matrix whose rows and columns are associated to basis pairs $e_i \wedge e_{i'}$ and $e_j \wedge e_{j'}$, where the ordering of the basis was $e_1 \wedge e_2, e_1 \wedge e_3$, and $e_2 \wedge e_3$. For instance, the upper right entry in the matrix for $\Lambda^2(L)$ is in its first row and third column, so it is the $e_1 \wedge e_2$ -coefficient of $\Lambda^2(L)(e_2 \wedge e_3)$. This matrix entry has row position $(1, 2)$ (index for the first basis vector) and column position $(2, 3)$ (index for the third basis vector). The 2×2 submatrix of L using rows 1 and 2 and columns 2 and 3 is $\begin{pmatrix} 2 & 0 \\ 1 & 1 \end{pmatrix}$, whose determinant is 2, which matches the upper right entry in the matrix at the end of Example 5.3.

7. EXTERIOR POWERS AND LINEAR INDEPENDENCE

This section discusses the connection between elementary wedge products and linear independence in a free module. We will start off with vector spaces, which are easier to handle.

If we are given n vectors in \mathbf{R}^n , there are two ways to determine if they are linearly independent using the $n \times n$ matrix with the vectors as the columns. The first way is to row reduce the matrix and see if you get the identity matrix. The second way is to compute the determinant of the matrix and see if you get a nonzero value. How can we decide if a set of $k < n$ vectors in \mathbf{R}^n is linearly dependent? Again there are two ways, each generalizing one of the previous two methods in terms of the $n \times k$ matrix having the vectors as the columns. The first way is to row reduce the matrix to see if a $k \times k$ submatrix is the identity. The second way is to compute the determinants of all $k \times k$ submatrices and see if any of them is not 0. Whereas the first way (row reduction) provides a set of steps that always keeps you going in the right direction, the second way involves one determinant computation after the other, and that will take longer to complete (particularly if all the determinants turn out to be 0!). Using exterior powers, we can carry out all these determinant computations at once. That is the algorithmic content of the next theorem for finite-dimensional vector spaces, if you keep in mind how determinants are related to coefficients in wedge product expansions.

Theorem 7.1. *Let V be a vector space. The vectors v_1, \dots, v_k in V are linearly independent if and only if $v_1 \wedge \dots \wedge v_k \neq 0$ in $\Lambda^k(V)$.*

Proof. The case $k = 1$ is trivial, so take $k \geq 2$. First assume $\{v_1, \dots, v_k\}$ is a linearly independent set. This set extends to a basis of V (we are working over a field!), so $v_1 \wedge \dots \wedge v_k$ is part of a basis of $\Lambda^k(V)$ by Theorem 4.2. In particular, $v_1 \wedge \dots \wedge v_k \neq 0$ in $\Lambda^k(V)$.

Now suppose $\{v_1, \dots, v_k\}$ is linearly dependent, so one v_i is a linear combination of the others. Whether or not $v_1 \wedge \dots \wedge v_k$ is nonzero in $\Lambda^k(V)$ is independent of the order of the factors, since permuting the terms only changes the elementary wedge product by a sign, so we may suppose v_k is a linear combination of the rest:

$$v_k = c_1 v_1 + \dots + c_{k-1} v_{k-1}.$$

Then in $\Lambda^k(V)$,

$$v_1 \wedge \dots \wedge v_{k-1} \wedge v_k = v_1 \wedge \dots \wedge v_{k-1} \wedge \left(\sum_{i=1}^{k-1} c_i v_i \right).$$

Expanding the right side gives a sum of $k - 1$ wedge products, each containing a repeated vector, so every term vanishes. \square

In this theorem V can be an infinite-dimensional vector space since we never required finite-dimensionality in the proof. At one point we invoked Theorem 4.2, which was proved for all free modules, not just free modules with a finite basis.

Corollary 7.2. *If V is a vector space and v_1, \dots, v_k are linearly independent in V , then an element $v \in V$ is a linear combination of v_1, \dots, v_k if and only if $v_1 \wedge \dots \wedge v_k \wedge v = 0$ in $\Lambda^{k+1}(V)$.*

Proof. Because the v_i 's are linearly independent, v is a linear combination of them if and only if $\{v_1, \dots, v_k, v\}$ is linearly dependent, and that is equivalent to their wedge product vanishing by Theorem 7.1. \square

Example 7.3. In \mathbf{R}^3 , does the vector $v = (4, -1, 1)$ lie in the span of $v_1 = (2, 1, 3)$ and $v_2 = (1, 2, 4)$? Since v_1 and v_2 are linearly independent, v is in their span if and only if $v_1 \wedge v_2 \wedge v$ vanishes in $\Lambda^3(\mathbf{R}^3)$. Let e_1, e_2, e_3 be the standard basis of \mathbf{R}^3 . Then we compute

$$\begin{aligned} v_1 \wedge v_2 &= (2e_1 + e_2 + 3e_3) \wedge (e_1 + 2e_2 + 4e_3) \\ &= 4(e_1 \wedge e_2) + 8(e_1 \wedge e_3) - e_1 \wedge e_2 + 4(e_2 \wedge e_3) \\ &\quad - 3(e_1 \wedge e_3) - 6(e_2 \wedge e_3) \\ &= 3(e_1 \wedge e_2) + 5(e_1 \wedge e_3) - 2(e_2 \wedge e_3) \end{aligned}$$

so

$$\begin{aligned} v_1 \wedge v_2 \wedge v &= (3(e_1 \wedge e_2) + 5(e_1 \wedge e_3) - 2(e_2 \wedge e_3)) \wedge (4e_1 - e_2 + e_3) \\ &= 3(e_1 \wedge e_2 \wedge e_3) - 5(e_1 \wedge e_3 \wedge e_2) - 8(e_2 \wedge e_3 \wedge e_1) \\ &= 3(e_1 \wedge e_2 \wedge e_3) + 5(e_1 \wedge e_2 \wedge e_3) - 8(e_1 \wedge e_2 \wedge e_3) \\ &= 0. \end{aligned}$$

Since this vanishes, $(4, -1, 1)$ is in the span of $(2, 1, 3)$ and $(1, 2, 4)$.

This method does not explicitly represent $(4, -1, 1)$ as a linear combination in the span of $(2, 1, 3)$ and $(1, 2, 4)$. (Here is one: $(4, -1, 1) = 3(2, 1, 3) - 2(1, 2, 4)$.) On the other hand, if we are only concerned with an existence question (is it in the span, not how is it in the

span) then this procedure works just as the computation of a determinant does to detect invertibility of a matrix without providing a formula for the inverse matrix.

There is a generalization of Theorem 7.1 to describe linear independence in a free R -module M rather than in a vector space. Given m_1, \dots, m_k in M , their linear independence is equivalent to a property of $m_1 \wedge \dots \wedge m_k$ in $\Lambda^k(M)$, but the property is *not* the nonvanishing of this elementary wedge product:

Theorem 7.4. *Let M be a free R -module. Elements m_1, \dots, m_k in M are linearly independent in M if and only if $m_1 \wedge \dots \wedge m_k$ in $\Lambda^k(M)$ is torsion-free: for $r \in R$, $r(m_1 \wedge \dots \wedge m_k) = 0$ only when $r = 0$.*

When R is a field, we recover Theorem 7.1. What makes Theorem 7.4 more subtle than Theorem 7.1 is that a linearly independent set in a free module usually does not extend to a basis, so we can't blindly adapt the proof of Theorem 7.1 to the case of free modules.

Proof. The case $k = 1$ is trivial by the definition of a linearly independent set in a module, so we can take $k \geq 2$.

Assume m_1, \dots, m_k is a linearly independent set in M . By Corollary 5.12, the elementary wedge product $m_1 \wedge \dots \wedge m_k$ is a linearly independent one-element subset in $\Lambda^k(M)$, so $m_1 \wedge \dots \wedge m_k$ has no R -torsion.

If the m_i 's are linearly dependent, say $r_1 m_1 + \dots + r_k m_k = 0$ with $r_i \in R$ not all 0, we may re-index and assume $r_1 \neq 0$. Then

$$0 = (r_1 m_1 + \dots + r_k m_k) \wedge m_2 \wedge \dots \wedge m_k = r_1 (m_1 \wedge m_2 \wedge \dots \wedge m_k),$$

so $m_1 \wedge m_2 \wedge \dots \wedge m_k$ has R -torsion. □

Corollary 7.5. *A system of d equations in d unknowns*

$$\begin{aligned} a_{11}x_1 + \dots + a_{1d}x_d &= 0 \\ a_{21}x_1 + \dots + a_{2d}x_d &= 0 \\ &\vdots \\ a_{d1}x_1 + \dots + a_{dd}x_d &= 0 \end{aligned}$$

in a commutative ring R has a nonzero solution $x_1, \dots, x_d \in R$ if and only if $\det(a_{ij})$ is a zero divisor in R .

Proof. We rewrite the theorem in terms of matrices: for $A \in M_d(R)$, the equation $Av = 0$ has a nonzero solution $v \in R^d$ if and only if $\det A$ is a zero divisor in R .

We consider the negated property, that the only solution of $Av = 0$ is $v = 0$. This is equivalent to the columns of A being linearly independent. The columns are Ae_1, \dots, Ae_d , and their elementary wedge product is

$$Ae_1 \wedge \dots \wedge Ae_d = (\det A)e_1 \wedge \dots \wedge e_d \in \Lambda^d(R^d).$$

Since $e_1 \wedge \dots \wedge e_d$ is a basis of $\Lambda^d(R^d)$ as an R -module, $(\det A)e_1 \wedge \dots \wedge e_d$ is torsion-free if and only if the only solution of $r \det A = 0$ is $r = 0$, which means $\det A$ is *not* a zero divisor. Thus $Av = 0$ has a nonzero solution if and only if $\det A$ is a zero divisor. □

We can turn Theorem 7.4 into a characterization of linear independence of k vectors v_1, \dots, v_k in R^d when $k \leq d$.³ The test is that $v_1 \wedge \dots \wedge v_k$ is torsion-free in $\Lambda^k(R^d)$. If we

³We may as well let $k \leq d$, since a linearly independent subset of R^d has at most d terms by Corollary 5.11.

let $A \in M_{d \times k}(R)$ be the matrix whose columns are v_1, \dots, v_k , linear independence of the columns is the same thing as injectivity of A as a linear map $R^k \rightarrow R^d$. So we can now say, for $k \leq d$, when a matrix in $M_{d \times k}(R)$ is injective as a linear transformation: if and only if the elementary wedge product of its columns is torsion-free in $\Lambda^k(R^d)$. What does that mean concretely? Writing $v_j = \sum_{i=1}^d a_{ij}e_i$ using the standard basis e_1, \dots, e_d of R^d ,

$$v_1 \wedge \cdots \wedge v_k = \sum_{1 \leq i_1 < \cdots < i_k \leq d} \begin{vmatrix} a_{1i_1} & \cdots & a_{1i_k} \\ \vdots & \ddots & \vdots \\ a_{ki_1} & \cdots & a_{ki_k} \end{vmatrix} e_{i_1} \wedge \cdots \wedge e_{i_k}.$$

This is torsion-free precisely when the coefficients are not all killed by a common nonzero element of R . Therefore we get the rule: for $k \leq d$, $A \in M_{d \times k}(R)$ is injective as a linear map $R^k \rightarrow R^d$ if and only if the determinants of its $k \times k$ submatrices have no common nonzero annihilator in R .

Example 7.6. Let

$$A = \begin{pmatrix} 2 & 2 \\ 1 & 5 \\ 1 & 2 \end{pmatrix} \in M_{3 \times 2}(\mathbf{Z}/6\mathbf{Z}).$$

Its 2×2 submatrices have determinants 8, -3 , 2, which equal 2, 3, 2 in $\mathbf{Z}/6\mathbf{Z}$. Although none of these determinants is a unit in $\mathbf{Z}/6\mathbf{Z}$, which would be an *easy* way to see injectivity, A is still injective because 2 and 3 have no common nonzero annihilator in $\mathbf{Z}/6\mathbf{Z}$ (that is, $2r = 0$ and $3r = 0$ only for $r = 0$). In the terminology of linear independence, we showed the two columns of A are linearly independent in $(\mathbf{Z}/6\mathbf{Z})^3$. (This does *not* mean neither column is a scalar multiple of the other, but it means the stronger assertion that no linear combination of the columns is 0 except for the trivial combination with both coefficients equal to 0.)

We can also check this in $\Lambda^2((\mathbf{Z}/6\mathbf{Z})^3)$: the two columns of A are $2e_1 + e_2 + e_3$ and $2e_1 + 5e_2 + 2e_3$, and

$$(2e_1 + e_2 + e_3) \wedge (2e_1 + 5e_2 + 2e_3) = 2e_1 \wedge e_2 + 2e_1 \wedge e_3 + 3e_2 \wedge e_3,$$

which is torsion-free since $2r = 0$ and $3r = 0$ in $\mathbf{Z}/6\mathbf{Z}$ only for $r = 0$.

That the coefficients in the elementary wedge product match the determinants of the 2×2 submatrices illustrates that the rules about elementary wedge product computations really do encode everything about determinants of submatrices.

I learned the following neat use of Corollary 7.5 from [2, pp. 6–7].

Theorem 7.7. *Let M be an R -module admitting a linear injection $R^d \hookrightarrow M$ and a linear surjection $R^d \twoheadrightarrow M$. Then $M \cong R^d$.*

Proof. The hypotheses say M has a d -element spanning set and a d -element linearly independent subset. Call the former x_1, \dots, x_d and the latter e_1, \dots, e_d , so we can write

$$M = Rx_1 + \cdots + Rx_d, \quad e_i = \sum_{j=1}^d a_{ij}x_j.$$

The x_j 's span M . We want to show they are linearly independent, so they form a basis and M is free of rank d .

The expression of the e_i 's in terms of the x_j 's can be written as the vector-matrix equation

$$(7.1) \quad \begin{pmatrix} e_1 \\ \vdots \\ e_d \end{pmatrix} = (a_{ij}) \begin{pmatrix} x_1 \\ \vdots \\ x_d \end{pmatrix},$$

where we treat $(a_{ij}) \in M_d(R)$ as a matrix acting on the d -tuples in M^d .

We show by contradiction that $\Delta := \det(a_{ij})$ is not a zero divisor in R . If Δ is a zero divisor then the columns of (a_{ij}) are linearly dependent by Corollary 7.5. A square matrix and its transpose have the same determinant, so the rows of (a_{ij}) are also linearly dependent. That means there are $c_1, \dots, c_d \in R$ not all 0 such that

$$(c_1, \dots, c_d)(a_{ij}) = (0, \dots, 0).$$

Using this, we multiply both sides of (7.1) on the left by (c_1, \dots, c_d) to get

$$(c_1, \dots, c_d) \begin{pmatrix} e_1 \\ \vdots \\ e_d \end{pmatrix} = 0.$$

That says $c_1e_1 + \dots + c_de_d = 0$, which contradicts linear independence of the e_i 's. So Δ is not a zero divisor.

Suppose now that $a_1x_1 + \dots + a_dx_d = 0$. We want to show every a_i is 0. Multiply both sides of (7.1) on the left by the cofactor matrix for (a_{ij}) (that's the matrix you multiply by to get the scalar diagonal matrix with the determinant Δ along the main diagonal):

$$\text{cof}(a_{ij}) \begin{pmatrix} e_1 \\ \vdots \\ e_d \end{pmatrix} = \Delta \begin{pmatrix} x_1 \\ \vdots \\ x_d \end{pmatrix}.$$

Now multiply both sides of this equation on the left by (a_1, \dots, a_d) :

$$(a_1, \dots, a_d) \text{cof}(a_{ij}) \begin{pmatrix} e_1 \\ \vdots \\ e_d \end{pmatrix} = \Delta(a_1x_1 + \dots + a_dx_d) = 0.$$

The product $(a_1, \dots, a_d) \text{cof}(a_{ij})$ is a (row) vector, say (b_1, \dots, b_d) . Then $b_1e_1 + \dots + b_de_d = 0$, so every b_i is 0 by linear independence of the e_i 's. Therefore $(a_1, \dots, a_d) \text{cof}(a_{ij}) = (0, \dots, 0)$. Multiply both sides of this on the *right* by (a_{ij}) to get $(a_1, \dots, a_d)\Delta = (0, \dots, 0)$, so $\Delta a_i = 0$ for all i . Since Δ is not a zero divisor, every a_i is 0 and we are done. \square

8. THE WEDGE PRODUCT

By concatenating elementary wedge products, we introduce a multiplication operation between different exterior powers of a module.

Lemma 8.1. *If $f: \underbrace{M \times \dots \times M}_{k \text{ times}} \times N_1 \times \dots \times N_\ell$ is multilinear and is alternating in the M 's,*

there is a unique multilinear map $\tilde{f}: \Lambda^k(M) \times N_1 \times \dots \times N_\ell \rightarrow N$ such that

$$\tilde{f}(m_1 \wedge \dots \wedge m_k, n_1, \dots, n_\ell) = f(m_1, \dots, m_k, n_1, \dots, n_\ell).$$

Proof. Uniqueness of \tilde{f} follows from multilinearity. To prove existence of \tilde{f} , fix a choice of $n_1 \in N_1, \dots, n_\ell \in N_\ell$. Define $f_{n_1, \dots, n_\ell}: M^k \rightarrow N$ by

$$f_{n_1, \dots, n_\ell}(m_1, \dots, m_k) = f(m_1, \dots, m_k, n_1, \dots, n_\ell).$$

Since f is multilinear, f_{n_1, \dots, n_ℓ} is multilinear. Since f is alternating in the m_i 's, f_{n_1, \dots, n_ℓ} is an alternating map. Therefore there is a unique linear map $\tilde{f}_{n_1, \dots, n_\ell}: \Lambda^k(M) \rightarrow N$ such that

$$\tilde{f}_{n_1, \dots, n_\ell}(m_1 \wedge \dots \wedge m_k) = f_{n_1, \dots, n_\ell}(m_1, \dots, m_k) = f(m_1, \dots, m_k, n_1, \dots, n_\ell).$$

on elementary wedge products.

Define $\tilde{f}: \Lambda^k(M) \times N_1 \times \dots \times N_\ell \rightarrow N$ by $\tilde{f}(\omega, n_1, \dots, n_\ell) = \tilde{f}_{n_1, \dots, n_\ell}(\omega)$. Since each f_{n_1, \dots, n_ℓ} is linear, \tilde{f} is linear in its first component. To check \tilde{f} is linear in one of its other coordinates, we carry it out for n_1 (all the rest are similar). We want to verify that

$$\tilde{f}(\omega, n_1 + n'_1, n_2, \dots, n_\ell) = \tilde{f}(\omega, n_1, n_2, \dots, n_\ell) + \tilde{f}(\omega, n'_1, n_2, \dots, n_\ell)$$

and

$$\tilde{f}(\omega, rn_1, n_2, \dots, n_\ell) = r\tilde{f}(\omega, n_1, n_2, \dots, n_\ell).$$

Both sides of both equations are additive in ω , so it suffices to check the identity when $\omega = m_1 \wedge \dots \wedge m_k$ is an elementary wedge product. In that case the two equations turn into linearity of f in its N_1 -component, and that is just a special case of the multilinearity of f . \square

Theorem 8.2. *Let M be an R -module. For positive integers k and ℓ , there is a unique R -bilinear map $\Lambda^k(M) \times \Lambda^\ell(M) \rightarrow \Lambda^{k+\ell}(M)$ satisfying the rule*

$$(m_1 \wedge \dots \wedge m_k, m'_1 \wedge \dots \wedge m'_\ell) \mapsto m_1 \wedge \dots \wedge m_k \wedge m'_1 \wedge \dots \wedge m'_\ell$$

on pairs of elementary wedge products.

Proof. Since the elementary wedge products span each exterior power module, and a bilinear map is determined by its values on pairs coming from spanning sets, there is *at most one* bilinear map with the prescribed behavior. As usual in this game, what needs proof is the *existence* of such a map.

Start by backing up and considering the function $f: M^k \times M^\ell \rightarrow \Lambda^{k+\ell}(M)$ by

$$f(m_1, \dots, m_k, m'_1, \dots, m'_\ell) = m_1 \wedge \dots \wedge m_k \wedge m'_1 \wedge \dots \wedge m'_\ell.$$

This is multilinear and alternating in the first k -coordinates, so by Lemma 8.1 there is a multilinear map $\tilde{f}: \Lambda^k(M) \times M^\ell \rightarrow \Lambda^{k+\ell}(M)$ such that

$$\tilde{f}(m_1 \wedge \dots \wedge m_k, m'_1, \dots, m'_\ell) = m_1 \wedge \dots \wedge m_k \wedge m'_1 \wedge \dots \wedge m'_\ell.$$

and it is alternating in its last ℓ coordinates. Therefore, again by Lemma 8.1, there is a bilinear map $B: \Lambda^k(M) \times \Lambda^\ell(M) \rightarrow \Lambda^{k+\ell}(M)$ such that

$$B(\omega, m'_1 \wedge \dots \wedge m'_\ell) = \tilde{f}(\omega, m'_1, \dots, m'_\ell),$$

so

$$\begin{aligned} B(m_1 \wedge \dots \wedge m_k, m'_1 \wedge \dots \wedge m'_\ell) &= \tilde{f}(m_1 \wedge \dots \wedge m_k, m'_1, \dots, m'_\ell) \\ &= m_1 \wedge \dots \wedge m_k \wedge m'_1 \wedge \dots \wedge m'_\ell. \end{aligned}$$

We have produced a bilinear map on the k th and ℓ th exterior powers with the desired value on pairs of elementary wedge products. \square

The operation constructed in Theorem 8.2, sending $\Lambda^k(M) \times \Lambda^\ell(M)$ to $\Lambda^{k+\ell}(M)$, is denoted \wedge and is called the *wedge product*. We place the operation in between the elements it acts on, just like other multiplication functions in mathematics. So for any $\omega \in \Lambda^k(M)$ and $\eta \in \Lambda^\ell(M)$ we have an element $\omega \wedge \eta \in \Lambda^{k+\ell}(M)$, and this operation is bilinear in ω and η . The formula in Theorem 8.2 on two elementary wedge products looks like this:

$$(m_1 \wedge \cdots \wedge m_k) \wedge (m'_1 \wedge \cdots \wedge m'_\ell) = m_1 \wedge \cdots \wedge m_k \wedge m'_1 \wedge \cdots \wedge m'_\ell.$$

Notice we have given a *new* meaning to the notation \wedge and to the terminology “wedge product,” which we have until now used in a purely formal way always in the phrase “elementary wedge product.” This new operational meaning of the wedge product is *consistent* with the old formal one, *e.g.* the (new) wedge product operation $\wedge: \Lambda^1(M) \times \Lambda^1(M) \rightarrow \Lambda^2(M)$ sends (m, m') to $m \wedge m'$ (old notation). This is very much like the definition of $R[T]$ as formal finite sums $\sum_i a_i T^i$ where, after the ring operations are defined, the symbol T^i is recognized as the i -fold product of the element T .

We have defined a wedge product $\Lambda^k(M) \times \Lambda^\ell(M) \rightarrow \Lambda^{k+\ell}(M)$ when k and ℓ are positive. What if one of them is 0? Recall (by definition) $\Lambda^0(M) = R$. Theorem 8.2 extends to the case $k = 0$ or $\ell = 0$ if we let the maps $\Lambda^0(M) \times \Lambda^\ell(M) \rightarrow \Lambda^\ell(M)$ and $\Lambda^k(M) \times \Lambda^0(M) \rightarrow \Lambda^k(M)$ be scalar multiplication: $r \wedge \eta = r\eta$ and $\omega \wedge r = r\omega$.

Now we can think about the old notation $m_1 \wedge \cdots \wedge m_k$ in an operational way: it is the result of applying the wedge product operation k times with elements from the module $\Lambda^1(M) = M$. Since we are now able to speak about a wedge product $\omega_1 \wedge \cdots \wedge \omega_k$ where the ω_i 's lie in exterior power modules $\Lambda^{k_i}(M)$, the elementary wedge products $m_1 \wedge \cdots \wedge m_k$, where the factors are in M , are merely special cases of wedge products.

Actually, to speak of $\omega_1 \wedge \cdots \wedge \omega_k$ unambiguously we need \wedge to be associative! So let's check that.

Theorem 8.3. *The wedge product is associative: if $\omega \in \Lambda^a(M)$, $\eta \in \Lambda^b(M)$, and $\xi \in \Lambda^c(M)$ then $(\omega \wedge \eta) \wedge \xi = \omega \wedge (\eta \wedge \xi)$ in $\Lambda^{a+b+c}(M)$.*

Proof. This is easy if a , b , or c is zero (then one of ω , η , and ξ is in R , and wedging with R is just scalar multiplication), so we can assume a , b , and c are all positive.

Let f and g be the functions from $\Lambda^a(M) \times \Lambda^b(M) \times \Lambda^c(M)$ to $\Lambda^{a+b+c}(M)$ given by both choices of parentheses:

$$f(\omega, \eta, \xi) = (\omega \wedge \eta) \wedge \xi, \quad g(\omega, \eta, \xi) = \omega \wedge (\eta \wedge \xi).$$

(Note $\omega \wedge \eta \in \Lambda^{a+b}(M)$ and $\eta \wedge \xi \in \Lambda^{b+c}(M)$.) Since the wedge product on two exterior power modules is bilinear, f and g are both trilinear functions. Therefore to show $f = g$ it suffices to verify equality on triples of elementary wedge products, since they are spanning sets of $\Lambda^a(M)$, $\Lambda^b(M)$, and $\Lambda^c(M)$. On such elements the equality is obvious by the definition of the wedge product operation, so we are done. \square

The following theorem puts associativity of the wedge product to work.

Theorem 8.4. *Let M be an R -module. If some $\Lambda^i(M)$ is finitely generated, then $\Lambda^j(M)$ is finitely generated for all $j \geq i$.*

The important point here is that we are not assuming M is finitely generated, only that some exterior power is finitely generated. Because exterior powers are not defined recursively, this theorem is not a tautology.⁴

Proof. It suffices to show that if $\Lambda^i(M)$ is finitely generated then $\Lambda^{i+1}(M)$ is finitely generated. Our argument is a simplification of [3, Lemma 2.1].

The module $\Lambda^i(M)$ has a finite spanning set, which we can take to be a set of elementary wedge products. Let $x_1, \dots, x_p \in M$ be the terms appearing in those elementary wedge products, so $\Lambda^i(M)$ is spanned by the i -fold wedges of x_1, \dots, x_p .

We will show $\Lambda^{i+1}(M)$ is spanned by the $(i+1)$ -fold wedges of x_1, \dots, x_p . It suffices to show the span of these wedges contains all the elementary wedge products in $\Lambda^{i+1}(M)$. Choose $(i+1)$ -fold elementary wedge product, say

$$y_1 \wedge \cdots \wedge y_i \wedge y_{i+1} = (y_1 \wedge \cdots \wedge y_i) \wedge y_{i+1}.$$

Since $y_1 \wedge \cdots \wedge y_i$ is in $\Lambda^i(M)$, it is an R -linear combination of i -fold wedges of x_1, \dots, x_p . Therefore $y_1 \wedge \cdots \wedge y_i \wedge y_{i+1}$ is an R -linear combination of expressions

$$(x_{j_1} \wedge \cdots \wedge x_{j_i}) \wedge y_{i+1} = x_{j_1} \wedge (x_{j_2} \wedge \cdots \wedge x_{j_i} \wedge y_{i+1}),$$

where the equation uses associativity of the wedge product. Since $x_{j_2} \wedge \cdots \wedge x_{j_i} \wedge y_{i+1}$ is in $\Lambda^i(M)$, it is an R -linear combination of i -fold wedges of x_1, \dots, x_p . \square

Theorem 8.5. *If M is spanned as an R -module by x_1, \dots, x_d then for $1 \leq k \leq d$ every element of $\Lambda^k(M)$ is a sum*

$$\omega_1 \wedge x_1 + \omega_2 \wedge x_2 + \cdots + \omega_d \wedge x_d$$

for some $\omega_i \in \Lambda^{k-1}(M)$.

We don't consider $k > d$ in this theorem since $\Lambda^k(M) = 0$ for such k by Theorem 4.1.

Proof. The result is clear when $k = 1$ since $\Lambda^0(M) = R$ by definition, so we can assume $k \geq 2$. From the beginning of Section 4, $\Lambda^k(M)$ is spanned as an R -module by the k -fold elementary wedge products $x_{i_1} \wedge \cdots \wedge x_{i_k}$ where $1 \leq i_1 < \cdots < i_k \leq d$. Using the wedge product multiplication $\Lambda^{k-1}(M) \times M \rightarrow \Lambda^k(M)$, we can write

$$x_{i_1} \wedge \cdots \wedge x_{i_k} = \eta \wedge x_{i_k},$$

where $\eta = x_{i_1} \wedge \cdots \wedge x_{i_{k-1}} \in \Lambda^{k-1}(M)$. Therefore every element of $\Lambda^k(M)$ is an R -linear combination $r_1(\eta_1 \wedge x_1) + \cdots + r_d(\eta_d \wedge x_d)$, where $r_i \in R$ and $\eta_i \in \Lambda^{k-1}(M)$. Since $r_i(\eta_i \wedge x_i) = (r_i \eta_i) \wedge x_i$ we can set $\omega_i = r_i \eta_i$ and we're done. \square

Here is an analogue of Theorem 4.8 for linear maps.

Theorem 8.6. *Let $\varphi: M \rightarrow N$ be a linear map of R -modules. If $\wedge^i(\varphi) = 0$ for some i then $\wedge^j(\varphi) = 0$ for all $j \geq i$.*

Theorem 4.8 is the special case when $N = M$ and $\varphi = \text{id}_M$ (why?).

⁴The theorem doesn't go backwards down to M , i.e., if some $\Lambda^i(M)$ is finitely generated this does not imply M is finitely generated. For example, $\Lambda^2(\mathbf{Q}/\mathbf{Z}) = 0$ as a \mathbf{Z} -module, but \mathbf{Q}/\mathbf{Z} is not finitely generated as a \mathbf{Z} -module.

Proof. It suffices so show $\wedge^{i+1}(\varphi) = 0$.

Since $\wedge^i(\varphi) = 0$ for any m_1, \dots, m_i in M we have

$$0 = \wedge^i(\varphi)(m_1 \wedge \cdots \wedge m_i) = \varphi(m_1) \wedge \cdots \wedge \varphi(m_i)$$

in $\Lambda^i(N)$. For any m_1, \dots, m_{i+1} in M , in $\Lambda^{i+1}(N)$ we have

$$\begin{aligned} \wedge^{i+1}(\varphi)(m_1 \wedge \cdots \wedge m_i \wedge m_{i+1}) &= \varphi(m_1) \wedge \cdots \wedge \varphi(m_i) \wedge \varphi(m_{i+1}) \\ &= (\varphi(m_1) \wedge \cdots \wedge \varphi(m_i)) \wedge \varphi(m_{i+1}) \\ &= 0 \wedge \varphi(m_{i+1}) \\ &= 0. \end{aligned}$$

Such terms span the image of $\wedge^{i+1}(\varphi)$, so $\wedge^{i+1}(\varphi) = 0$. \square

In an elementary wedge product $m_1 \wedge \cdots \wedge m_k$, where the factors are in $M = \Lambda^1(M)$, transposing two of them introduces a sign change. What is the sign-change rule for transposing factors in a wedge product $\omega_1 \wedge \cdots \wedge \omega_k$? By associativity, we just need to understand how a single wedge product $\omega \wedge \eta$ changes when the factors are reversed.

Theorem 8.7. For $\omega \in \Lambda^k(M)$ and $\eta \in \Lambda^\ell(M)$, $\omega \wedge \eta = (-1)^{k\ell} \eta \wedge \omega$.

Proof. This is trivial if $k = 0$ or $\ell = 0$ (in which case the wedge product is simply scaling and $(-1)^{k\ell} = 1$), so we can take k and ℓ positive.

Both sides of the desired equation are bilinear functions of ω and η , so to verify equality for all ω and η it suffices to do so on spanning sets of the modules. Thus we can take $\omega = m_1 \wedge \cdots \wedge m_k$ and $\eta = m'_1 \wedge \cdots \wedge m'_\ell$, so we want to show

$$(8.1) \quad m_1 \wedge \cdots \wedge m_k \wedge m'_1 \wedge \cdots \wedge m'_\ell = (-1)^{k\ell} (m'_1 \wedge \cdots \wedge m'_\ell \wedge m_1 \wedge \cdots \wedge m_k).$$

Starting with the expression on the left, we successively move $m'_1, m'_2, \dots, m'_\ell$ to the front. First move m'_1 past each of the m_i 's, which is a total of k swaps, so

$$m_1 \wedge \cdots \wedge m_k \wedge m'_1 \wedge \cdots \wedge m'_\ell = (-1)^k m'_1 \wedge m_1 \wedge \cdots \wedge m_k \wedge m'_2 \wedge \cdots \wedge m'_\ell.$$

Now move m'_2 past every m_i , introducing another set of k sign changes:

$$m_1 \wedge \cdots \wedge m_k \wedge m'_1 \wedge \cdots \wedge m'_\ell = (-1)^{2k} m'_1 \wedge m'_2 \wedge m_1 \wedge \cdots \wedge m_k \wedge m'_3 \wedge \cdots \wedge m'_\ell.$$

Repeat until m'_ℓ has been moved past every m_i . In all, there are $k\ell$ swaps, so the overall sign at the end is $(-1)^{k\ell}$. \square

Theorem 8.8. For odd k and $\omega \in \Lambda^k(M)$, $\omega \wedge \omega = 0$.

Proof. There is a quick proof using Theorem 8.7 when $2 \in R^\times$: $\omega \wedge \omega = (-1)^{k^2} (\omega \wedge \omega) = -(\omega \wedge \omega)$, so $2(\omega \wedge \omega) = 0$, so $\omega \wedge \omega = 0$.

To handle the general case when 2 may not be a unit, write

$$\omega = \sum_{i_1, \dots, i_k} c_{i_1, \dots, i_k} m_{i_1} \wedge \cdots \wedge m_{i_k} = \sum_I c_I \omega_I,$$

where $I = (i_1, \dots, i_k)$ is a multi-index of k integers and $\omega_I = \omega_{(i_1, \dots, i_k)}$ is an abbreviation for $m_{i_1} \wedge \cdots \wedge m_{i_k}$. Then, by the bilinearity of the wedge product,

$$(8.2) \quad \omega \wedge \omega = \sum_I c_I m_I \wedge \sum_J c_J \omega_J = \sum_{I, J} c_I c_J \omega_I \wedge \omega_J,$$

where I and J run over the same set of multi-indices. Note each $\omega_I \wedge \omega_J$ is an elementary wedge product with $2k$ factors.

The multi-indices I and J could be equal. In that case $\omega_I \wedge \omega_J = 0$ since it is a $2k$ -fold elementary wedge product with repeated factors. When $I \neq J$ the double sum in (8.2) contains

$$c_I c_J \omega_I \wedge \omega_J + c_J c_I \omega_J \wedge \omega_I = c_I c_J (\omega_I \wedge \omega_J + \omega_J \wedge \omega_I).$$

Since ω_I and ω_J are in $\Lambda^k(M)$, $\omega_J \wedge \omega_I = (-1)^{k^2} (\omega_I \wedge \omega_J) = -(\omega_I \wedge \omega_J)$, so $c_I c_J (\omega_I \wedge \omega_J + \omega_J \wedge \omega_I) = 0$. \square

What about Theorem 8.8 when k is even? If $\omega = m_1 \wedge \cdots \wedge m_k$ is an *elementary* wedge product in $\Lambda^k(M)$ then $\omega \wedge \omega$ vanishes since it is an elementary wedge product with a repeated factor from M . But it is not generally true that $\omega \wedge \omega = 0$ for all $\omega \in \Lambda^k(M)$.

Example 8.9. For $k \geq 2$, let M be finite free with linearly independent subset e_1, \dots, e_{2k} . Set $\omega = e_1 \wedge \cdots \wedge e_k + e_{k+1} \wedge \cdots \wedge e_{2k} \in \Lambda^k(M)$. This is a sum of two elementary wedge products, and

$$\omega \wedge \omega = 2e_1 \wedge \cdots \wedge e_k \wedge e_{k+1} \wedge \cdots \wedge e_{2k} \in \Lambda^{2k}(M).$$

By Corollary 5.12 and Theorem 7.4, $e_1 \wedge \cdots \wedge e_k \wedge e_{k+1} \wedge \cdots \wedge e_{2k}$ is torsion-free in $\Lambda^{2k}(M)$, so when $2 \neq 0$ in R we have $\omega \wedge \omega \neq 0$. Elementary wedge products always “square” to 0, so ω is *not* an elementary wedge product when $2 \neq 0$ in R .

To get practice computing in an exterior power module, we look at the equation $v \wedge \omega = 0$ in $\Lambda^{k+1}(V)$, where V is a vector space and $\omega \in \Lambda^k(V)$.

Theorem 8.10. *Let V be a vector space. For nonzero $\omega \in \Lambda^k(V)$,*

$$\dim(\{v \in V : v \wedge \omega = 0\}) \leq k,$$

with equality if and only if ω is an elementary wedge product.

Proof. The result is obvious if $k = 0$, so we can suppose $k \geq 1$. Let v_1, \dots, v_d be linearly independent vectors in V that each wedge ω to 0. We want to show $d \leq k$. Since V might be infinite-dimensional, we first create a suitable finite-dimensional subspace W in which we can work. (If you want to assume V is finite-dimensional, set $W = V$ and skip the rest of this paragraph.) Let W be the span of v_1, \dots, v_d and the nonzero vectors appearing in some fixed representation of ω as a finite sum of elementary wedge products in $\Lambda^k(V)$. Then W is finite-dimensional. There’s a natural embedding $W \hookrightarrow V$ and we get an embedding $\Lambda^\ell(W) \hookrightarrow \Lambda^\ell(V)$ in a natural way for all ℓ by Corollary 5.9. If we view each v_i in W and ω in $\Lambda^k(W)$ then the condition $v_i \wedge \omega = 0$ in $\Lambda^{k+1}(V)$ implies $v_i \wedge \omega = 0$ in $\Lambda^{k+1}(W)$.

Set $n = \dim W$, so obviously $d \leq n$. If $k \geq n$ then obviously $d \leq k$, so we can suppose $k \leq n - 1$.

Extend $\{v_1, \dots, v_d\}$ to a basis $\{v_1, \dots, v_n\}$ of W . Using this basis we can write ω in $\Lambda^k(W)$ as a finite sum of linearly independent elementary wedge products:

$$\omega = \sum_{1 \leq i_1 < \cdots < i_k \leq n} c_{i_1, \dots, i_k} v_{i_1} \wedge \cdots \wedge v_{i_k}$$

and some coefficient is not 0. Fix i between 1 and d , and compute $v_i \wedge \omega$ using this formula:

$$\begin{aligned} 0 &= v_i \wedge \omega \\ &= \sum_{1 \leq i_1 < \dots < i_k \leq n} c_{i_1, \dots, i_k} v_i \wedge v_{i_1} \wedge \dots \wedge v_{i_k} \\ &= \sum_{\substack{1 \leq i_1 < \dots < i_k \leq n \\ i_1, \dots, i_k \neq i}} c_{i_1, \dots, i_k} v_i \wedge v_{i_1} \wedge \dots \wedge v_{i_k}. \end{aligned}$$

This equation is taking place in $\Lambda^{k+1}(W)$, where the $(k+1)$ -fold wedges $v_i \wedge v_{i_1} \wedge \dots \wedge v_{i_k}$ for $i \notin \{i_1, \dots, i_k\}$ are linearly independent. Therefore the coefficients here are all 0:

$$i \notin \{i_1, \dots, i_k\} \Rightarrow c_{i_1, \dots, i_k} = 0.$$

Here i was any number from 1 to d , so

$$c_{i_1, \dots, i_k} \neq 0 \Rightarrow \{1, \dots, d\} \subset \{i_1, \dots, i_k\}.$$

There is at least one nonzero coefficient, so we must have $\{1, \dots, d\} \subset \{i_1, \dots, i_k\}$ for some k -tuple of indices. Counting the two sets, $d \leq k$.

If $d = k$ then $\{1, \dots, k\} = \{i_1, \dots, i_k\}$, which allows just one nonzero term and $\omega = c_{1, \dots, k} v_1 \wedge \dots \wedge v_k$, which is an elementary wedge product. Conversely, if ω is a nonzero elementary wedge product in $\Lambda^k(V)$ then $\{v \in V : \omega \wedge v = 0\}$ has dimension k by Corollary 7.2 and associativity of the wedge product. \square

Theorem 8.11. *Let V be a vector space and $k \geq 1$. For nonzero $v \in V$ and $\omega \in \Lambda^k(V)$, $v \wedge \omega = 0$ if and only if $\omega = v \wedge \eta$ for some $\eta \in \Lambda^{k-1}(V)$.*

Proof. By associativity, if $\omega = v \wedge \eta$ then $v \wedge \omega = v \wedge (v \wedge \eta) = (v \wedge v) \wedge \eta = 0$. The point of the theorem is that the converse direction holds: if $v \wedge \omega = 0$ then we can write $\omega = v \wedge \eta$ for some η .

As in the proof of the previous theorem, we can reduce to the finite-dimensional case (details left to the reader), so we'll just take V to be a finite-dimensional vector space. Set $n = \dim V \geq 1$. Since $\omega \in \Lambda^k(V)$, if $k > n$ then $\omega = 0$ and we can trivially write $\omega = v \wedge 0$. So we may suppose $k \leq n$. If $k = 0$ then $\omega \in \Lambda^0(V)$ is a scalar and wedging with ω is scalar multiplication, so the conditions $v \neq 0$ and $v \wedge \omega = 0$ imply $\omega = 0$. Thus again $\omega = v \wedge 0$.

Now suppose $1 \leq k \leq n$. Extend v to a basis of V , say v_1, \dots, v_n where $v = v_1$. If $k = n$ the condition $v \wedge \omega = 0$ is automatic since $\Lambda^{n+1}(V) = 0$. And it is also automatic that ω is “divisible” by v if $k = n$: $\Lambda^n(V)$ has basis $v_1 \wedge \dots \wedge v_n$, so $\omega = c(v_1 \wedge \dots \wedge v_n)$ for some scalar c . Then $\omega = v \wedge \eta$ where $\eta = cv_2 \wedge \dots \wedge v_n$.

We now assume $1 \leq k \leq n - 1$ (so $n \geq 2$). Using the basis of $\Lambda^k(V)$ coming from our chosen basis of V that includes v as the first member v_1 ,

$$(8.3) \quad \omega = \sum_{1 \leq i_1 < \dots < i_k \leq n} c_{i_1, \dots, i_k} v_{i_1} \wedge \dots \wedge v_{i_k}$$

with scalar coefficients. Since $v_1 = v$, $v \wedge v_1 = 0$ so

$$(8.4) \quad v \wedge \omega = \sum_{2 \leq i_1 < \dots < i_k \leq n} c_{i_1, \dots, i_k} v \wedge v_{i_1} \wedge \dots \wedge v_{i_k}.$$

The elementary wedge products on the right are part of a basis of $\Lambda^{k+1}(V)$, so from $v \wedge \omega = 0$ we see all the coefficients in (8.4) vanish. Thus in (8.3) the only nonzero terms are among

those with $i_1 = 1$, so we can pull out $v_1 = v$:

$$\omega = v \wedge \sum_{1 < i_2 < \dots < i_k \leq n} c_{1, i_2, \dots, i_k} v_{i_2} \wedge \dots \wedge v_{i_k}.$$

Let η be the large sum here, so $\omega = v \wedge \eta$. □

9. THE EXTERIOR ALGEBRA

Since wedge products move elements into higher-degree exterior powers, we can view \wedge as multiplication in a noncommutative ring by taking the direct sum of all the exterior powers of a module.

Definition 9.1. For an R -module M , its *exterior algebra* is the direct sum

$$\Lambda(M) = \bigoplus_{k \geq 0} \Lambda^k(M) = R \oplus M \oplus \Lambda^2(M) \oplus \Lambda^3(M) \oplus \dots,$$

provided with the multiplication rule given by the wedge product from Theorem 8.2, extended distributively to the whole direct sum.

Each $\Lambda^k(M)$ is only an R -module, but their direct sum $\Lambda(M)$ is an R -algebra: it has a multiplication which commutes with scaling by R and has identity $(1, 0, 0, \dots)$. Notice R is a subring of $\Lambda(M)$, embedded as $r \mapsto (r, 0, 0, \dots)$.

When $\Lambda^k(M)$ is viewed in the exterior algebra $\Lambda(M)$, it is called a homogeneous part and its elements are said to be the homogeneous terms of *degree* k in $\Lambda(M)$. For instance, when m_0, m_1, m_2, m_3 , and m_4 are in M the sum $m_0 + m_1 \wedge m_2 + m_3 \wedge m_4$ in $\Lambda(M)$ has homogeneous parts m_0 and $m_1 \wedge m_2 + m_3 \wedge m_4$. This is the same terminology used for homogeneous multivariable polynomials, e.g., $X^2 - XY + Y^3$ has homogeneous parts $X^2 - XY$ (of degree 2) and Y^3 (of degree 3).

The wedge product on the exterior algebra $\Lambda(M)$ is bilinear, associative, and distributive, but not commutative (unless $-1 = 1$ in R). The replacement for commutativity in Theorem 8.7 does *not* generalize to a rule between all elements of $\Lambda(M)$. However, if $\omega \in \Lambda^k(M)$ and k is even then ω commutes with every element of $\Lambda(M)$, so more generally the submodule $\bigoplus_{k \text{ even}} \Lambda^k(M)$ lies in the center of $\Lambda(M)$.

A typical element of $\Lambda(M)$ is a sequence $(\omega_k)_{k \geq 0}$ with $\omega_k = 0$ for $k \gg 0$, and we write it as a formal sum $\sum_{k \geq 0} \omega_k$ while keeping the direct sum aspect in mind. In this notation, the wedge product of two elements of $\Lambda(M)$ is

$$\sum_{k \geq 0} \omega_k \wedge \sum_{\ell \geq 0} \eta_\ell = \sum_{p \geq 0} \left(\sum_{k+\ell=p} \omega_k \wedge \eta_\ell \right),$$

where the inner sum on the right is actual addition in $\Lambda^p(M)$ and the outer sum is purely formal (corresponding to the direct sum decomposition defining $\Lambda(M)$). This extension of the wedge product from operations $\Lambda^k(M) \times \Lambda^\ell(M) \rightarrow \Lambda^{k+\ell}(M)$ on different exterior power modules to a single operation $\Lambda(M) \times \Lambda(M) \rightarrow \Lambda(M)$ is analogous to the way the multiplication rule $(aT^i)(bT^j) = abT^{i+j}$ on monomials, which is associative, can be extended to the usual multiplication between any two polynomials in $R[T] = \bigoplus_{i \geq 0} RT^i$. When M is a finitely generated R -module with n generators, every $\Lambda^k(M)$ is finitely generated as an R -module and $\Lambda^k(M) = 0$ for $k > n$ so $\Lambda(M) = \bigoplus_{k=0}^n \Lambda^k(M)$ is finitely generated as an R -module. Because $\Lambda^k(M)$ is spanned as an R -module by the elementary wedge products,

and an elementary wedge product is a wedge product of elements of M , $\Lambda(M)$ is generated as an R -algebra (not as an R -module!) by M .

Theorem 9.2. *When M is a free R -module of rank d , its exterior algebra is a free R -module of rank 2^d .*

Proof. Each exterior power module $\Lambda^k(M)$ is free with rank $\binom{d}{k}$ for $0 \leq k \leq d$ and vanishes for $k > d$, so their direct sum $\Lambda(M)$ is free with rank

$$\sum_{k=0}^d \binom{d}{k} = 2^d.$$

□

Concretely, when M is free with basis e_1, \dots, e_d , we can think of $\Lambda(M)$ as an R -algebra generated by the e_i 's subject to the relations $e_i^2 = 0$ and $e_i e_j = -e_j e_i$ for $i \neq j$. The construction of $\Lambda(M)$ was basis-free, but this explicit description when there is a basis is helpful when doing computations.

Example 9.3. Let V be a real vector space of dimension 3 with basis e_1, e_2, e_3 . Then $\Lambda^0(V) = \mathbf{R}$, $\Lambda^1(V) = V = \mathbf{R}e_1 \oplus \mathbf{R}e_2 \oplus \mathbf{R}e_3$, $\Lambda^2(V) = \mathbf{R}(e_1 \wedge e_2) \oplus \mathbf{R}(e_1 \wedge e_3) \oplus \mathbf{R}(e_2 \wedge e_3)$, and $\Lambda^3(V) = \mathbf{R}(e_1 \wedge e_2 \wedge e_3)$. The exterior algebra $\Lambda(V)$ is the direct sum of these vector spaces and we can count the dimension as $1 + 3 + 3 + 1 = 8$.

To get practice computing in an exterior algebra, we ask which elements of the exterior algebra of a vector space wedge a given vector to 0.

Theorem 9.4. *Let V be a vector space over the field K . For nonzero $v \in V$ and $\omega \in \Lambda(V)$, $v \wedge \omega = 0$ if and only if $\omega = v \wedge \eta$ for some $\eta \in \Lambda(V)$.*

Proof. The reduction to the case of finite-dimensional V proceeds as in the reduction step of the proof of Theorem 8.11. By associativity, if $\omega = v \wedge \eta$ then $v \wedge \omega = v \wedge (v \wedge \eta) = (v \wedge v) \wedge \eta = 0$. The point of the theorem is that the converse direction holds. First we reduce to the finite-dimensional case. Let W be the span of v and all the nonzero elementary wedge products in an expression for ω . Since $W \hookrightarrow V$, $\Lambda^k(W) \hookrightarrow \Lambda^k(V)$, so $\Lambda(W) \hookrightarrow \Lambda(V)$. From these embeddings, it suffices to prove the theorem with V replaced by W , so we may assume V is finite-dimensional.

Write $\omega = \sum_{k=0}^n \omega_k$ where $\omega_k \in \Lambda^k(V)$ is the degree k part of ω . (If you think about direct sums as sequences, $\omega = (\omega_0, \dots, \omega_n)$.) Then

$$v \wedge \omega = \sum_{k=0}^n v \wedge \omega_k.$$

Since $v \wedge \omega_k \in \Lambda^{k+1}(V)$, the terms in the sum are in different homogeneous parts of $\Lambda(V)$, so the vanishing of $v \wedge \omega$ implies $v \wedge \omega_k = 0$ for each k . By Theorem 8.11, $\omega_0 = 0$ and for $k \geq 1$ we have $\omega_k = v \wedge \eta_{k-1}$ for some $\eta_{k-1} \in \Lambda^{k-1}(V)$. Thus $\omega = v \wedge \eta$ where $\eta = \sum_{k=1}^n \eta_{k-1} \in \Lambda(V)$. □

While we created the exterior algebra $\Lambda(M)$ as a direct sum of R -modules with a snazzy multiplicative structure, it can be characterized on its own terms among R -algebras by a universal mapping property. Since each $\Lambda^k(M)$ is spanned as an R -module by the elementary wedge products $m_1 \wedge \dots \wedge m_k$, $\Lambda(M)$ is generated as an R -algebra (using wedge

multiplication) by M . Moreover, $m \wedge m = 0$ in $\Lambda(M)$ for all $m \in M$. We now turn this into a universal mapping property.

Theorem 9.5. *Let A be any R -algebra and suppose there is an R -linear map $L: M \rightarrow A$ such that $L(m)^2 = 0$ for all $m \in M$. Then there is a unique extension of L to an R -algebra map $\tilde{L}: \Lambda(M) \rightarrow A$. That is, there is a unique R -algebra map \tilde{L} making the diagram*

$$\begin{array}{ccc} \Lambda(M) & & \\ \uparrow & \searrow \tilde{L} & \\ M & \xrightarrow{L} & A \end{array}$$

commute.

This theorem is saying that any R -linear map from M to an R -algebra A such that the image elements square to 0 can always be extended uniquely to an R -algebra map from $\Lambda(M)$ to A . Such a universal property determines $\Lambda(M)$ up to R -algebra isomorphism by the usual argument.

Proof. Since M generates $\Lambda(M)$ as an R -algebra, there is at most one R -algebra map $\Lambda(M) \rightarrow A$ whose values on M are given by L . The whole problem is to construct such a map.

For any m and m' in M , $L(m)^2 = 0$, $L(m')^2 = 0$, and $L(m + m')^2 = 0$. Expanding $(L(m) + L(m'))^2$ and removing the squared terms leaves

$$0 = L(m)L(m') + L(m')L(m),$$

which shows $L(m)L(m') = -L(m')L(m)$. Therefore any product of L -values on M can be permuted at the cost of an overall sign change. This implies $L(m_1)L(m_2) \cdots L(m_k) = 0$ if two m_i 's are equal, since we can permute the terms to bring them together and then use the vanishing of $L(m_i)^2$. This will be used later.

If there is going to be an R -algebra map $f: \Lambda(M) \rightarrow A$ extending L , then on an elementary wedge product we must have

$$f(m_1 \wedge \cdots \wedge m_k) = f(m_1) \cdots f(m_k) = L(m_1) \cdots L(m_k).$$

since \wedge is the multiplication in $\Lambda(M)$. To show there is such a map, we start on the level of the $\Lambda^k(M)$'s. For $k \geq 0$, let $M^k \rightarrow A$ by $(m_1, \dots, m_k) \mapsto L(m_1) \cdots L(m_k)$. This is multilinear since L is R -linear. It is alternating because $L(m_1) \cdots L(m_k) = 0$ when two m_i 's are equal. Hence we obtain an R -linear map $f_k: \Lambda^k(M) \rightarrow A$ satisfying

$$f_k(m_1 \wedge \cdots \wedge m_k) = L(m_1) \cdots L(m_k)$$

for all elementary wedge products $m_1 \wedge \cdots \wedge m_k$. Define $f: \Lambda(M) \rightarrow A$ by letting it be f_k on $\Lambda^k(M)$ and extending to the direct sum $\Lambda(M) = \bigoplus_{k \geq 0} \Lambda^k(M)$ by additivity. This function f is R -linear and it is left to the reader to show f is multiplicative. \square

Notice the individual $\Lambda^k(M)$'s don't appear in the statement of Theorem 9.5. This theorem describes an intrinsic feature of the full exterior algebra as an R -algebra.

REFERENCES

- [1] W. C. Brown, "A Second Course in Linear Algebra," J. Wiley & Sons, New York, 1988.
- [2] H. Flanders, *Tensor and Exterior Powers*, J. Algebra **7**, 1–24 (1967).
- [3] R. Gardner, *Modules Admitting Determinants*, Linear and Multilinear Algebra **3** (1975/76), 209–214.
- [4] J. Rotman, "Advanced Modern Algebra," Prentice-Hall, Upper Saddle River, NJ, 2002.
- [5] F. Warner, "Foundations of Differentiable Manifolds and Lie Groups," Springer-Verlag, New York, 1983.