

MATH 395. BASES OF SYMMETRIC AND EXTERIOR POWERS

Let V be a finite-dimensional nonzero vector spaces over a field F , say with dimension d . For any $n \geq 1$, the n th symmetric and exterior powers $\text{Sym}^n(V)$ and $\wedge^n(V)$ were made as quotients of $V^{\otimes n}$ that “universally linearize” symmetric and alternating multilinear mappings $V^{\times n} \rightarrow W$. Our aim here is to find bases of these spaces in terms of bases of V . We fix an ordered basis $\mathbf{e} = \{e_1, \dots, e_d\}$ of V .

1. PRELIMINARY CONSIDERATIONS

Let $\mu : V^{\times n} \rightarrow W$ be a multilinear mapping. For any $v_1, \dots, v_n \in V$, say $v_j = \sum_{i=1}^d a_{ij}e_i$, multilinearity gives

$$(1) \quad \mu(v_1, \dots, v_n) = \sum_{i_1, \dots, i_n} (a_{i_1,1} \cdots a_{i_n,n}) \mu(e_{i_1}, \dots, e_{i_n})$$

with $1 \leq i_1, \dots, i_n \leq d$. Conversely, for arbitrary $w_{i_1, \dots, i_n} \in W$ we can *define*

$$\mu(v_1, \dots, v_n) = \sum_{i_1, \dots, i_n} (a_{i_1,1} \cdots a_{i_n,n}) w_{i_1, \dots, i_n}$$

for $v_j = \sum_{i=1}^d a_{ij}e_i$ to get a multilinear mapping (check!) satisfying $\mu(e_{i_1}, \dots, e_{i_n}) = w_{i_1, \dots, i_n}$. In other words, to give a multilinear μ is “the same” as to give vectors $w_{i_1, \dots, i_n} \in W$ indexed by ordered n -tuples of integers between 1 and $d = \dim V$. This correspondence depends on \mathbf{e} , and it is a restatement of the fact that the n^d elementary tensors $e_{i_1} \otimes \cdots \otimes e_{i_n}$ are a basis of $V^{\otimes n}$. Hence, it is natural to expect that properties of μ may be read off from these n^d values of μ . We claim that if the values of μ on the n -tuples from the basis satisfy the symmetry or skew-symmetry conditions, then the same holds for μ in general. That is:

Lemma 1.1. *The multilinear μ is symmetric if and only if the value $\mu(e_{i_1}, \dots, e_{i_n})$ is always invariant under switching i_j and $i_{j'}$ for any distinct $1 \leq j, j' \leq n$, and μ is skew-symmetric if this value always negates upon switching i_j and $i_{j'}$ for any distinct $1 \leq j, j' \leq n$.*

Taking $V = F^2$ and $\mu((x, y), (x', y')) = xy'$ gives an example with $\mu(e_1, e_1) = \mu(e_2, e_2) = 0$ (for $e_1 = (1, 0)$ and $e_2 = (0, 1)$) but μ is not alternating. Indeed, $\mu(e_1 + e_2, e_1 + e_2) = 1$. Thus, the alternating property cannot be detected just from looking at the values of μ on vectors from a basis.

Proof. In each case, the necessity of the asserted property of the behavior of the values $\mu(e_{i_1}, \dots, e_{i_n})$ is a special case of the general property for μ to be symmetric or skew-symmetric. For the converse, look at (1). If we switch v_k and $v_{k'}$ then we simply switch a_{ik} and $a_{ik'}$. Thus, the coefficient of the term for (i_1, \dots, i_n) is modified by replacing $a_{i_k, k}$ and $a_{i_{k'}, k'}$ with $a_{i_k, k'}$ and $a_{i_{k'}, k}$ respectively. This is, if $k < k'$, then the coefficients against $\mu(e_{i_1}, \dots, e_{i_n})$ and $\mu(e_{i_1}, \dots, e_{i_{k'}}, \dots, e_{i_k}, \dots, e_{i_n})$ are swapped. But these μ -values are equal (resp. negatives of each other) under the hypothesis of symmetry (resp. skew-symmetry) for μ on ordered n -tuples of basis vectors! Hence, under such a condition of symmetry (resp. skew-symmetry) for μ on ordered n -tuples from the basis, we deduce the property of symmetry (resp. skew-symmetry) for μ in general. ■

In general, suppose $\mu : V^{\times n} \rightarrow W$ is a multilinear mapping that is symmetric (resp. alternating). By (1), the multilinearity ensures that to uniquely determine μ , we just need to specify the values

$$\mu(e_{i_1}, \dots, e_{i_n}) \in W$$

for $1 \leq i_1, \dots, i_n \leq d$, and that these may be specified *arbitrarily*. If μ is to be symmetric, then this list has redundancies: for any ordered n -tuple of i_j 's, the assigned value must equal the one

assigned to the order n -tuple obtained by rearranging the i_j 's in monotonically increasing order. That is, for symmetric μ we only need to restrict attention to specifying values as above in the special case $1 \leq i_1 \leq \dots \leq i_n \leq d$.

If μ is to be alternating, or more generally skew-symmetric, then we may again rearrange the i_j 's to be in monotonically increasing order, say via some permutation σ of the j 's (i.e., we replace i_j with $i_{\sigma(j)}$ for some $\sigma \in \mathfrak{S}_n$). The value of μ prior to shuffling around the e_{i_j} 's is related to the value after the i_j 's are arranged in monotonically increasing order via a factor of the sign of σ . If the i_j 's are pairwise distinct then this permutation σ is uniquely determined (as there is a unique way to shuffle the j 's to put the i_j 's in strictly increasing order). For alternating μ , we lose nothing by restricting attention to the case of pairwise distinct i_j 's, as in all other cases the value of μ has to be zero. Thus, for alternating μ we only need to specify the values

$$\mu(e_{i_1}, \dots, e_{i_n})$$

for $1 \leq i_1 < \dots < i_n \leq d$. (In particular, if $n > d$ then no such strictly increasing sequence exists, which is to say that there is *always* a repetition of some i_j 's, and hence necessarily $\mu = 0$: that is, if $n > \dim V$ then an alternating mapping $\mu : V^{\times n} \rightarrow W$ is automatically zero.)

We are now motivated to ask if, upon specifying the values $\mu(e_{i_1}, \dots, e_{i_n}) \in W$ for all monotone increasing (resp. strictly increasing) sequences of i_j 's between 1 and d , there actually exists a symmetric (resp. alternating) $\mu : V^{\times n} \rightarrow W$ realizing these specified values. The preceding shows that such a μ is unique, and in the language of symmetric and exterior powers of V the problem is precisely that of determining if a linear map $\text{Sym}^n(V) \rightarrow W$ (resp. $\wedge^n(V) \rightarrow W$) may be constructed by arbitrarily specifying its values on n -fold products $e_{i_1} \cdots e_{i_n}$ (resp. $e_{i_1} \wedge \cdots \wedge e_{i_n}$) of basis vectors e_{i_j} with the i_j 's monotone increasing (resp. strictly increasing). In other words, do such elementary products provide a basis of the n th symmetric (resp. exterior) power of V ? The answer is yes:

Theorem 1.2. *For any $n \geq 1$, a basis of $\text{Sym}^n(V)$ is given by the n -fold products $e_{i_1} \cdots e_{i_n}$ for $1 \leq i_1 \leq \dots \leq i_n \leq d$. For $n > d$ the space $\wedge^n(V)$ vanishes and if $1 \leq n \leq d$ then $\wedge^n(V)$ has a basis given by the n -fold products $e_{i_1} \wedge \cdots \wedge e_{i_n}$ for $1 \leq i_1 < \dots < i_n \leq d$.*

The vanishing of $\wedge^n(V)$ for $n > \dim V$ has been shown above: for such n , any alternating multilinear map $\mu : V^{\times n} \rightarrow W$ has been proved to be zero, and so it uniquely factors through the vanishing map to the zero vector space. Hence, the unique characterization of the exterior power via its mapping property thereby ensures $\wedge^n(V) = 0$ in such cases. Thus, in the proof of Theorem 1.2 for the case of exterior powers we may (and do) restrict attention to the case $1 \leq n \leq d$.

2. PROOF OF THEOREM 1.2

To see the spanning aspect, one can use the universal properties of symmetric and exterior powers much as in our analogous argument for why tensors of basis vectors span a tensor product. However, we can also give a direct proof as follows. By construction, the symmetric and exterior powers are quotients of tensor powers, and under this quotient map an elementary tensor of an ordered n -tuple of vectors in V is mapped to the corresponding elementary n -fold symmetric or wedge product of the same ordered set of vectors. Thus, since spanning sets of a vector space map to spanning sets in a quotient space, we conclude that n -fold symmetric (resp. wedge) products of e_i 's span the n th symmetric (resp. exterior) power of V since their elementary n -fold tensors span $V^{\otimes n}$. This gives the spanning result using e_{i_j} 's without any ordering restriction on the i_j 's. However, the symmetry (resp. skew-symmetry) of the mapping $V^{\times n} \rightarrow \text{Sym}^n(V)$ (resp. $V^{\times n} \rightarrow \wedge^n(V)$) allows us to rearrange the i_j 's in monotone-increasing order at the possible expense of some signs (which

is harmless for the purposes of being a spanning set). Thus, we get the spanning result using just $i_1 \leq \dots \leq i_n$ in the case of symmetric powers, and for exterior powers we get the spanning result using $i_1 < \dots < i_n$ because this only ignores the cases when there is a repetition amongst the i_j 's and in such cases the wedge product of the e_{i_j} 's vanishes. This concludes the proof that the asserted bases for the symmetric and exterior powers of V are at least spanning sets.

Now we have to prove linear independence. We imitate the method for tensor products: the aim is to construct linear functionals on the symmetric and exterior powers that kill all but exactly one of the elements of the proposed basis, with this distinguished element allowed to be specified arbitrarily in advance. Applying such functionals to any potential linear relation would force the corresponding coefficients in the linear relation to vanish, and hence all coefficients would have to vanish (as desired).

First we handle the alternating case (as this turns out to be slightly easier). Suppose $n \leq d$ and fix an ordered n -tuple $I = (i_1, \dots, i_n)$ with $1 \leq i_1 < \dots < i_n \leq d$. Define $B_I : V^{\times n} \rightarrow F$ to be the multilinear form whose value on (v_1, \dots, v_n) is the determinant of the $n \times n$ submatrix using rows i_1, \dots, i_n in the $d \times n$ matrix formed by the v_j 's viewed in F^d via \mathbf{e} -coordinates. That is, if $v_j = \sum_{i=1}^d a_{ij} e_i$ then define

$$\mu_I(v_1, \dots, v_n) = \det(a_{i_r, j})_{1 \leq r, j \leq n} = \sum_{\sigma \in \mathfrak{S}_n} \text{sign}(\sigma) \prod_{j=1}^n e_{i_{\sigma(j)}}^*(v_j)$$

where e_1^*, \dots, e_d^* is the dual basis to e_1, \dots, e_d . By the definition, one checks that μ_I is indeed multilinear and alternating. Thus, it uniquely factors through a linear functional $T_I : \wedge^n(V) \rightarrow F$ with

$$T_I(v_1 \wedge \dots \wedge v_n) = \mu_I(v_1, \dots, v_n).$$

If $I' = (i'_1, \dots, i'_n)$ is an ordered n -tuple with $1 \leq i'_1 < \dots < i'_n \leq d$ then the strict monotonicity implies that when $I' \neq I$ some i'_{j_0} must be distinct from all of the i_j 's and hence $\mu_I(e_{i'_1}, \dots, e_{i'_n})$ vanishes (because the relevant submatrix must have an entire j_0 th column of zeros, or alternatively each product in the big summation formula has a factor of zero for $j = j_0$). Thus,

$$T_I(e_{i'_1} \wedge \dots \wedge e_{i'_n}) = 0$$

when $I' \neq I$, and clearly the value is 1 when $I' = I$ (the determinant is for the ‘‘identity matrix’’). If there is a linear relation

$$\sum_{i'_1 < \dots < i'_n} c_{i'_1, \dots, i'_n} e_{i'_1} \wedge \dots \wedge e_{i'_n} = 0$$

in $\wedge^n(V)$ then applying the *linear* functional T_I kills all but precisely the I th term, giving

$$c_{i_1, \dots, i_n} = 0.$$

Varying I , we get that all coefficients vanish. This is the desired linear independence in $\wedge^n(V)$.

Now we turn to the case of symmetric powers. In this case there is an analogue of μ_I but some factorials arise when there is a repetition among the i_j 's and so in order to handle the possibility of fields F with positive characteristic (in which some factorials may vanish) we have to be a little more cunning in our definition of the analogue of μ_I . Here is an initial attempt at an analogous construction. For $I = (i_1, \dots, i_n)$ with $1 \leq i_1 \leq \dots \leq i_n \leq d$ and any $v_j = \sum_{i=1}^d a_{ij} e_i \in V$ ($1 \leq j \leq n$) we define

$$\mu'_I(v_1, \dots, v_n) = \sum_{\sigma \in \mathfrak{S}_n} \prod_{j=1}^n e_{i_{\sigma(j)}}^*(v_j);$$

this is like μ_I except that we have removed the signs. The removal of the signs makes this expression multilinear and symmetric (check!). If we let $T'_I : \text{Sym}^n(V) \rightarrow F$ be the resulting linear functional, so

$$T'_I(v_1 \cdots v_n) = \mu'_I(v_1, \dots, v_n),$$

then when $I' = (i'_1, \dots, i'_n)$ is a monotonically increasing sequence with $I' \neq I$ we must have some i'_{j_0} that is distinct from all $i_j \in I$ (as a monotone sequence can be made monotone in exactly one way), so $T'_I(e_{i'_1} \cdots e_{i'_n}) = 0$ when $I' \neq I$.

This is adequate to complete the proof of linear independence, much as in the alternating case, provided that

$$T_I(e_{i_1} \cdots e_{i_n}) \neq 0$$

in F . We shall now compute this value, and we will see that when F has positive characteristic it may vanish in F . Suppose there are $N \leq n$ distinct values i_j as j varies (so $N = n$ if there are no repetitions, which is to say that $i_j \neq i_{j'}$ whenever $j \neq j'$). Consider the partitioning of the set $\{1, \dots, n\}$ into pairwise disjoint subsets J_1, \dots, J_N of indices j for which the i_j 's have a common value. (By the monotonicity of the i_j 's, this partitioning has elements of J_r less than those of J_s whenever $r < s$.) If there are no repetitions then $J_r = \{r\}$. Note that N and the J_r 's depend on I , but I is fixed for now.

Let $n_r = \#J_r$ denote the size of J_r . There are $n_r!$ permutations of the elements of J_r , and as we let σ vary over all $\prod_{r=1}^N (n_r!)$ permutations of $\{1, \dots, n\}$ that permute each of the J_r 's, the rearrangements of I again return I . Consequently, the sum

$$T'_I(e_{i_1} \cdots e_{i_n}) = \mu(e_{i_1}, \dots, e_{i_n})$$

has exactly $\prod_{r=1}^N (n_r!)$ terms equal to 1 and all others equal to 0. Thus, the value is this product of factorials considered in F . Hence, we are done if F has characteristic 0 or characteristic $> n$, but otherwise for some I one of these factorials will vanish in F and so T'_I will vanish. Hence, to get the linear independence result without restriction on the characteristic of F we need to build a symmetric multilinear form that works better than μ'_I .

The trick is to introduce the subgroup $H \subseteq \mathfrak{S}_n$ consisting of those permutations σ of $\{1, \dots, n\}$ such that $\sigma(J_r) = J_r$ for all r ; that is, σ permutes the J_r 's. Note that H depends on our fixed choice of I (as the J_r 's do). The elements of H are exactly the σ 's such that the ordered n -tuple $(i_{\sigma(1)}, \dots, i_{\sigma(n)})$ is equal to I , since $\sigma \in H$ precisely when $h(m) \in J_r$ if and only if $m \in J_r$, which is to say (by the definition of the J_r 's) that $i_{h(m)} = i_m$ for all $1 \leq m \leq n$. In particular, H may be identified with the product of the permutation groups of the J_r 's, and so H has size $\prod_{r=1}^N (n_r!)$.

For any $\sigma \in \mathfrak{S}_n$ and $v_1, \dots, v_n \in V$, the product

$$\prod_{j=1}^n e_{i_{\sigma(j)}}^*(v_j) \in F$$

only depends on the left coset $H\sigma$ since for any $h \in H$ we have $i_{(h\sigma)(j)} = i_{h(\sigma(j))} = i_{\sigma(j)}$. Writing $H \backslash \mathfrak{S}_n$ to denote the set of left H -cosets, to an element $\bar{\sigma} \in H \backslash \mathfrak{S}_n$ we may associate the product

$$\prod_{j=1}^n e_{i_{\sigma(j)}}^*(v_j) \in F$$

using *any* representative $\sigma \in \mathfrak{S}_n$ for the coset $\bar{\sigma}$. Thus, we consider the modified mapping $M_I : V^{\times n} \rightarrow F$ defined by

$$M_I(v_1, \dots, v_n) = \sum_{\bar{\sigma} \in H \backslash \mathfrak{S}_n} \prod_{j=1}^n e_{i_{\sigma(j)}}^*(v_j) \in F.$$

Each product involves a linear functional evaluated on each of the v_j 's exactly once, and so M_I is a multilinear mapping. Moreover, it is readily check that it is symmetric (essentially because \mathfrak{S}_n has a well-defined *right* multiplication action on the set of left cosets $H \backslash \mathfrak{S}_n$).

Roughly speaking, M_I improves μ'_I by eliminating the “ H -fold” repetition in the sum defining μ'_I . The exact same argument as for T'_I shows that the linear functional ℓ_I on $\text{Sym}^n(V)$ induced by M_I kills the symmetric product $e_{i'_1} \cdots e_{i'_n}$ whenever $I' = (i'_1, \dots, i'_n)$ is a monotone ordered n -tuple distinct from I (as each individual term of the sum defining $M_I(e_{i'_1}, \dots, e_{i'_n})$ vanishes). But now in the case $I' = I$ we win because the value on $e_{i_1} \cdots e_{i_n}$ is

$$M_I(e_{i_1}, \dots, e_{i_n}) = 1.$$

The point is that cutting down on the number of terms in the sum by means of left H -cosets exactly removes the multiplicity of $\#H = \prod_{r=1}^N (n_r!)$ in the earlier calculation, and so gives the value of 1.

The multilinear symmetric M_I 's provide linear functionals ℓ_I on $\text{Sym}^n(V)$ that permit us to prove the desired linear independence exactly as in the alternating case: each ℓ_I kills $e_{i'_1} \cdots e_{i'_n}$ whenever the monotone ordered n -tuple $I' = (i'_1, \dots, i'_n)$ is distinct from I , and it has the value 1 in the case $I' = I$. Thus, applying the ℓ_I 's to any potential linear dependence relation among these elementary symmetric products forces each coefficient in such a relation in $\text{Sym}^n(V)$ to vanish.