1 Multilinear and alternating functions

We have already encountered multilinear polynomials in the proofs of intersection theorems of Frankl–Wilson and Ray-Chaudhuri–Wilson. We will now extend the notion to functions on an arbitrary product of vector spaces.

Let $W_1, \dots, W_k, V$ be vector spaces over some field $\mathbb{F}$. A function

$$f\colon W_1 \times \cdots \times W_k \to V$$

is *multilinear* if it is linear in each of the k variables separately. In other words,

$$f(\dots, w_{i-1}, \sum_j \alpha_j w_{ij}, w_{i+1}, \dots) = \sum_j \alpha_j f(\dots, w_{i-1}, w_{ij}, w_{i+1}, \dots),$$

for any choice of $w_1 \in W_1$, $w_2 \in W_2$, etc.

There are special names for multilinear functions when k is small. When $k = 2$ we speak of a bilinear function functions, and of trilinear functions when $k = 3$. In general, we may call f a *k-linear* function.

Example: Recall that a polynomial on $\mathbb{F}^n$ is called multilinear if it is a linear combination of monomials of the form $x_I \stackrel{\text{def}}{=} \prod_{i \in I} x_i$. Such a polynomial is the same as a multilinear function $f\colon \mathbb{F} \times \cdots \times \mathbb{F} \to \mathbb{F}$.

A special class of multilinear functions are alternating functions. If f is a multilinear function, $W_1 = \ldots = W_k$ and $f(w_1, \dots, w_k) = 0$ whenever $w_i = w_j$ for some $i \neq j$, then we say that f is *alternating*.

Example: Given a vectors $w_1, \dots, w_n \in \mathbb{F}_n$ let $[w_1, \dots, w_n]^T$ be the n-by-n matrix whose rows are $w_1, \dots, w_n$ in that order. Then the function $D(w_1, \dots, w_n) = \det[w_1, \dots, w_n]^T$ is alternating.

Lemma 1. *Let $f\colon W \times \cdots \times W \to V$ be an alternating function. Then the following hold:*

a) $f(\dots, w_{i-1}, \sum \alpha_j w_j, w_{i+1}, \dots) = \alpha_i f(\dots, w_{i-1}, w_i, w_{i+1}, \dots)$

b) $f(\dots, w_i, \dots, w_j, \dots) = -f(\dots, w_j, \dots, w_i, \dots)$.

c) If $e_1, \dots, e_n$ are a basis for W, then for every $w_1, \dots, w_k \in W$ we have

$$f(w_1, \dots, w_k) \in \operatorname{span}\{f(e_{j_1}, \dots, e_{j_k}) : j_1 < \cdots < j_k\}.$$

The property (b) is the one responsible for the name 'alternating'. If $\operatorname{char}\mathbb{F} \neq 2$, then it is easy to see that a multilinear functions satisfying property (b) is alternating in our sense. However, if $\operatorname{char}\mathbb{F} = 2$, the implication fails, as witnessed by the function $\mathbb{F}_2 \times \mathbb{F}_2 \to \mathbb{F}_2$ that is non-zero only when both of its arguments are non-zero.

Proof of Lemma 1.

a) This follows by invoking multilinearity, and noting all the terms save for one vanish.

b) This is a consequence of (a) and the following computation:

$$\begin{aligned} f(\ldots, w_i, \ldots, w_j, \ldots) &= f(\ldots, w_i + w_j, \ldots, w_j, \ldots) \\ &= f(\ldots, w_i + w_j, \ldots, w_j - (w_i + w_j), \ldots) \\ &= -f(\ldots, w_i + w_j, \ldots, w_i, \ldots) \\ &= -f(\ldots, w_j, \ldots, w_i, \ldots). \end{aligned}$$

c) Let $w_i = \sum_j a_{ij} e_j$. Then by multilinearity

$$f(w_1, \ldots, w_n) = \sum_{j_1, \ldots, j_k} \alpha_{1j_1} \cdot \cdots \cdot \alpha_{kj_k} f(e_{j_1}, \ldots, e_{j_k}).$$

The terms in which same basis vectors appears twice vanish, whereas the remaining terms can be arranged so that $j_1 < j_2 < \ldots < j_k$ by repeated application of (b).

□

2 Exterior powers

In this section we shall work with vector space $W = \mathbb{F}^n$ exclusively. The part (c) of Lemma 1 tells us that if f is an alternating function on W^k, then $f(W^k)$ spans an at most $\binom{n}{k}$-dimensional space. Intuitively, $\dim \operatorname{span} f(W^k) = \binom{n}{k}$ holds whenever there are no relations between values of f other than those implied by the alternating property of f. We shall make this intuition rigorous below, but for now we exhibit an alternating function such that $f(W_k)$ spans a $\binom{n}{k}$-dimensional space.

Theorem 2. *For each k, there exists a $\binom{n}{k}$-dimensional space T_k and an alternating function $f_k \colon \mathbb{F}^n \times \cdots \times \mathbb{F}^n \to T_k$, and, for each k and l, there exists a bilinear function $g_{kl} \colon T_k \times T_l \to T_{k+l}$ that satisfy the following two properties:*

a) $\dim \operatorname{span}\{f_k(w_1,\dots,w_k) : w_1,\dots,w_k \in \mathbb{F}^n\} = \binom{n}{k}$,

b) $g_{kl}\big(f_k(w_1,\dots,w_k), f_l(u_1,\dots,u_l)\big) = f_{k+l}(w_1,\dots,w_k,u_1,\dots,u_l)$.

Proof. We define $T_k \stackrel{\text{def}}{=} \mathbb{F}^{\binom{[n]}{k}}$. In other words, T_k a $\binom{n}{k}$-tuples of elements from $\mathbb{F}$ indexed by k-element subsets of $[n]$.

Given vectors $w_1,\dots,w_k \in \mathbb{F}^n$, form a k-by-n matrix A whose rows are $w_1,\dots,w_k$ (in that order). Then for a set $I \in \binom{[n]}{k}$ let A_I be the k-by-k submatrix of A that is made of columns indexed by the set I. Then define $f_k \colon W \times \cdots \times W \to T_k$ by

$$f_k(w_1,\dots,w_k)_I = \det A_I.$$

The function f is multinear because det is multilinear, and it is alternating because det is alternating. To see that f_k satisfies part (a), consider any k distinct basis vectors $e_{i_1},\dots,e_{i_k}$ in $\mathbb{F}^n$. Then

$$f(e_{i_1},\dots,e_{i_k})_I = \begin{cases} \pm 1 & \text{if } I = \{i_1,\dots,i_k\}, \\ 0 & \text{otherwise,} \end{cases}$$

where the sign is determined by the order of $i_1,\dots,i_k$. So, each element of the standard basis of T_k is in the image of f, and so span $f(W^k) = T_k$.

Defining g_{kl} is simple but cumbersome: we first let $\Sigma(I) \stackrel{\text{def}}{=} \sum_{i\in I} i$, then for each $I \in \binom{[n]}{k+l}$ we put

$$g_{kl}(x,y)_I = (-1)^{\Sigma([k])} \sum_{\substack{I_1\cup I_2 = I \\ I_1 \in \binom{[n]}{k} \\ I_2 \in \binom{[n]}{l}}} (-1)^{\Sigma(I_1)} x_{I_1} y_{I_2}.$$

The part (b) of the lemma amounts to the identity

$$\det C = (-1)^{\Sigma([k])} \sum_{\substack{I_1\cup I_2 = I \\ I_1 \in \binom{[n]}{k} \\ I_2 \in \binom{[n]}{l}}} (-1)^{\Sigma(I_1)} \det A_{I_1} \cdot \det B_{I_2}$$

where C is a matrix of the form $C = \left[\begin{smallmatrix} A \\ B \end{smallmatrix}\right]$ where A, B are the k-by-n, and l-by-n matrices respectively. To see the identity, note that both sides are polynomials made of the same monomials, and we just need to check that the signs match. The latter is verified by noting that the signs match for $I_1 = [k]$, and that swapping an integer from I_1 with one in I_2 changes the sign correctly. $\square$

The vector spaces T_k are called *exterior powers* of $\mathbb{F}^n$, and are traditionally denoted by $\bigwedge^k \mathbb{F}^n$. The functions f_k and g_{kl} are called *alternating product* and denoted by the same wedge symbol $\wedge$. In particular, the identity $g_{kl}(f_k, f_l) = f_{k+l}$ from part (b) of the preceding lemma is written concisely as

$$(w_1 \wedge \cdots \wedge w_k) \wedge (v_1 \wedge \cdots \wedge v_l) = w_1 \wedge \cdots \wedge w_k \wedge v_1 \wedge \cdots \wedge v_l,$$

and expresses associativity of the alternating product. Of course, the alternating product is not commutative since $w \wedge w' = -w' \wedge w$.

3 From subspaces to vectors

The k-dimensional subspaces in $\mathbb{F}^n$ can be naturally regarded as vectors in $\bigwedge^k \mathbb{F}^n$. Indeed, given a subspace U with a basis $u_1, \ldots, u_k$ we can consider the vector

$$\wedge U \stackrel{\text{def}}{=} u_1 \wedge \cdots \wedge u_k.$$

The vector is non-zero since the matrix $A = [u_1, \ldots, u_k]^T$ has rank k, and so $\det A_I \neq 0$ for some k-by-k minor.

As defined, $\wedge U$ depends on the choice of the basis in U, but it does so only slightly. If $u'_1, \ldots, u'_k$ is any other basis, then by part (c) of Lemma 1 it follows that $u'_1 \wedge \cdots \wedge u'_k = \lambda u_1 \wedge \cdots \wedge u_k$ for some scalar λ. Hence, we can regard $\wedge U$ as a vector in the projective space $(\wedge^k \mathbb{F}^n)/\mathbb{F}^*$. Below we will abuse notation and pretend that $\wedge U$ is an element of $\wedge^k \mathbb{F}^n$.

Note that $(\wedge U) \wedge v = 0$ if and only if $v \in \wedge U$. In particular, this implies that $\wedge U$ completely determines the set of vectors in U, i.e., the map $U \mapsto \wedge U$ is injective. However, not every vector in $\bigwedge^k \mathbb{F}^n$ is of the form $\wedge U$. Indeed, the space of all k-dimensional subspaces of $\mathbb{F}^n$ is $k(n-k)$-dimensional[1], whereas $\dim \bigwedge \mathbb{F}^n = \binom{n}{k}$ is much larger.

A special case important in the computer graphics is the case $k = 2$ and $n = 4$. Thinking projectively, a 2-dimensional subspace of $\mathbb{R}^4$ corresponds to a line in $\mathbb{R}^3$. Thus, the map $U \mapsto \wedge U$ associates to each line in $\mathbb{R}^3$ a sixtuple of numbers. This sixtuple is known as *Plücker coordinates* of the line.

[1]The collection of all k-dimensional subspaces forms an algebraic variety, called the *Grassmannian*. As we have not yet defined the notion of dimension for algebraic varieties, the statement is informal: to specify a k-dimensional subspace of $\mathbb{F}^n$ we 'need' $k(n-k)$ scalars.

4 Bollobás-type theorems

We start by giving another proof of the skew version of Bollobás's theorem. The proof is due to Lovász.

Theorem 3. *Suppose $(A_1, B_1), \dots, (A_m, B_m)$ are pairs of sets that satisfy*

a) A_i is an r-element set, and B_i is an s-element set, for all i,

b) $A_i \cap B_i = \emptyset$, for all i,

c) $A_i \cap B_j \neq \emptyset$, for all $i < j$.

Then the number of pairs is $m \le \binom{r+s}{s}$.

Proof. Let X be a finite set containing all the A_i's and B_i's. Let $W = \mathbb{R}^{r+s}$ and associate to each $x \in X$ a vector w_x such that the vectors $\{w_x : x \in X\}$ are in general position, i.e., no $r+s$ of them are linearly dependent.

To each subset $S \subset X$ associate a vector $w_S \overset{\text{def}}{=} \bigwedge_{x \in S} w_s$. With this association we have, for any sets $A, B \subset X$ satisfying $|A| + |B| \le r + s$,

$$w_A \cap w_B = 0 \iff A \cap B \neq \emptyset.$$

Indeed, if $A \cap B \neq \emptyset$, then the alternating product $w_A \wedge w_B$ contains a repeated element. Conversely, if $A \cap B = \emptyset$, then $A \cup B$ is a linearly independent set, and so $w_A \wedge w_B$ is non-zero.

In particular,

$$w_{A_i} \wedge w_{B_i} \begin{cases} \neq 0 & \text{if } i = j, \\ = 0 & \text{if } i < j, \end{cases}$$

and so the vectors $w_{A_1}, \dots, w_{A_m} \in \bigwedge^r \mathbb{R}^{r+s}$ are linearly independent. Since $\dim \bigwedge^k \mathbb{R}^{r+s} = \binom{r+s}{r}$, it follows that $m \le \binom{r+s}{r}$. $\square$

In Bollobás's theorem the pairs are classified according to whether they intersect or not. In the following extension of the result, due to Füredi, the pairs are classified by the size of the intersection.

Theorem 4. *Let t be a nonnegative integer. Suppose $(A_1, B_1), \dots, (A_m, B_m)$ are pairs of sets that satisfy*

a) A_i is an r-element set, and B_i is an s-element set, for all i,

b) $|A_i \cap B_i| \le t$, for all i,

c) $|A_i \cap B_j| \ge t + 1$, for all $i < j$.

Then the number of pairs is $m \le \binom{r+s-2t}{r-t}$.

The result is tight. To see that, set the ground set to be $[r+s-2t]$, and let A_i be a union of $[t]$ and any $(r-t)$-element subset of $[r+s-t] \setminus [t]$, and B_i be the union of $[t]$ and remaining elements from $[r+s-t] \setminus [t]$.

We will deduce 4 from a related result for subspace intersections.

Theorem 5. *Suppose* $(U_1, V_1), \ldots, (U_m, V_m)$ *are pairs of subspaces of a vectors space over a field* $\mathbb{F}$. *Assume that*

a) U_i *is* r*-dimensional, and* B_i *is* s*-dimensional, for all* i*,*

b) $\dim(A_i \cap B_i) \le t$*, for all* i*,*

c) $\dim(A_i \cap B_j) \ge t+1$*, for all* $i < j$.

Then $m \le \binom{r+s-2t}{r-t}$.

Lemma 6. *Suppose* $U_1, \ldots, U_r$ *are subspaces of a finite-dimensional vector space* W *over an infinite field* $\mathbb{F}$*, and let* d *be a nonnegative integer not exceeding* $\dim W$. *Then the following hold:*

a) There exists a linear mapping $\phi\colon W \to \mathbb{F}^d$ *such that*

$$\dim \phi(U_i) = \min(\dim U_i, d).$$

b) There exists a codimension d *subspace* L *such that*

$$\dim(L \cap U_i) = \max(\dim U_i - d, 0).$$

We postpone the proof of these assertion until our discussion of algebraic varieties later in the course. Intuitively, since $\mathbb{F}$ is infinite, 'almost all' maps ϕ (resp. subspaces L) satisfy any one of these conditions, and there are only finitely many of these conditions. One can also prove this assertion directly working in coordinates (exercise!). An easier exercise is to show that the two assertions are equivalent.

Proof of 5. We may assume that the field $\mathbb{F}$ is infinite; if $\mathbb{F}$ is finite, we can replace $\mathbb{F}$ by its algebraic closure.

By part (b) of Lemma 6 there exists a subspace L of codimension t that satisfies

$$\begin{aligned} \dim(U_i \cap L) &= r-t, \\ \dim(V_i \cap L) &= s-t, \\ \dim(U_i \cap V_i \cap L) &= 0, \\ \dim(U_i \cap V_j \cap L) &\ge 1, \qquad \text{for } i < j. \end{aligned}$$

Let $U_i' \stackrel{\text{def}}{=} U_i \cap L$ and $V_i' \stackrel{\text{def}}{=} V_i \cap L$. We have thus reduced the problem to the case $t = 0$ for the pairs $(U_1', V_1'), \dots, (U_m', V_m') \in \mathbb{F}^{r-t} \times \mathbb{F}^{s-t}$. The next step is to reduce the dimension of the ambient space.

By part (a) of Lemma 6 there exists a linear map $\phi \colon W \to \mathbb{F}^{r+s-2t}$ such that the spaces $\tilde{U}_i \stackrel{\text{def}}{=} \phi(U_i')$ and $\tilde{V}_j \stackrel{\text{def}}{=} \phi(V_j')$ have the same dimensions as U_i' and V_j' respectively, and $\dim \operatorname{span} \tilde{U}_i \cup \tilde{V}_j = \dim \operatorname{span} U_i' \cup V_j'$. The latter condition implies $U_i' \cap V_j' = \{0\}$ if and only if $\tilde{U}_i \cap \tilde{V}_j = \{0\}$.

Let $u_i = \wedge \tilde{U}_i$ and $v_j = \wedge \tilde{V}_j$ It follows that

$$u_i \wedge v_j \begin{cases} \neq 0 & \text{if } i = j, \\ = 0 & \text{if } i < j, \end{cases}$$

and so the vectors $u_1, \dots, u_m \in \bigwedge^r W$ are linearly independent. Since these are vectors in $\bigwedge^r \mathbb{F}^{r+s-2t}$, we conclude that $m \leq \binom{r+s-2t}{r-t}$. $\square$

Proof of Theorem 4. Without loss the ground set is $[n]$. Let $e_1, \dots, e_n$ be the standard basis vectors for $\mathbb{F}^n$. Associate to each set S the vectors space $V_S = \operatorname{span}\{e_i : i \in S\}$. Then the vector space pairs $(V_{A_1}, V_{B_1}), \dots, (V_{A_m}, V_{B_m})$ satisfy the assumption of Theorem 5. $\square$