Math 395. Tensor products and bases

Let V and V' be finite-dimensional vector spaces over a field F. Recall that a *tensor product* of V and V' is a pait (T, t) consisting of a vector space T over F and a bilinear pairing $t : V \times V' \to T$ with the following universal property: for *any* bilinear pairing $B : V \times V' \to W$ to *any* vector space W over F, there exists a unique linear map $L : T \to W$ such that $B = L \circ t$. Roughly speaking, t "uniquely linearizes" all bilinear pairings of V and V' into arbitrary F-vector spaces.

In class it was proved that if (T, t) and (T', t') are two tensor products of V and V', then there exists a unique linear isomorphism $T \simeq T'$ carrying t and t' (and *vice-versa*). In this sense, the tensor product of V and V' (equipped with its "universal" bilinear pairing from $V \times V'$!) is unique up to *unique* isomorphism, and so we may speak of "the" tensor product of V and V'. You must never forget to think about the data of t when you contemplate the tensor product of V and V': it is the pair (T, t) and not merely the underlying vector space T that is the focus of interest. In this handout, we review a method of construction of tensor products (there is another method that involved no choices, but is horribly "big"-looking and is needed when considering modules over commutative rings) and we work out some examples related to the construction.

As was indicated in class, we shall write $V \times V' \stackrel{\otimes}{\to} V \otimes V'$ to denote a tensor product (if one exists), and in particular for $v \in V$ and $v' \in V'$ the image of (v, v') under the "universal" bilinear pairing into $V \otimes V'$ shall be denoted $v \otimes v' \in V \otimes V'$. By virtue of the bilinearity of the pairing $(v, v') \mapsto v \otimes v'$, we have relations such as

$$(a_1 v_1 + a_2 v_2) \otimes v' = a_1 (v_1 \otimes v') + a_2 (v_2 \otimes v')$$

in $V \otimes V'$ for any $a_1, a_2 \in F$, $v_1, v_2 \in V$, and $v' \in V'$.

Note that if $V = 0$ or $V' = 0$, then $V \otimes V' = 0$. Indeed, the only bilinear pairing $B : V \times V' \to W$ is the zero pairing, and hence the pairing $V \times V' \to \{0\}$ given by the zero pairing certainly fits the bill to be the tensor product. We therefore say $V \otimes 0 = 0$ and $0 \otimes V' = 0$. This example is not particularly interesting, but anyway it permits us to focus the existence problem on the more interesting case when $n = \dim V$ and $n' = \dim V'$ are both positive.

1. Some generalities with bilinear pairings

The key to everything is the following elementary lemma.

Lemma 1.1. *Let $\{v_i\}_{1 \le i \le n}$ and $\{v'_j\}_{1 \le j \le n'}$ be ordered bases of V and V'. To give a bilinear pairing $B : V \times V' \to W$ is "the same" as to give nn' vectors $w_{ij} = B(v_i, v'_j) \in W$ in the following sense: any B is uniquely determined by the pairings $B(v_i, v'_j) \in W$, and conversely if $w_{ij} \in W$ are arbitrarily given then there exists a (necessarily unique!) bilinear pairing $B : V \times V' \to W$ satisfying $B(v_i, v'_j) = w_{ij}$.*

Proof. Suppose we are given a bilinear pairing $B : V \times V' \to W$. Any $v \in V$ and $v' \in V'$ admit unique expansions $v = \sum a_i v_i$ and $v' = \sum b'_j v'_j$, so by bilinearity we have

$$B(v, v') = \sum_{i,j} a_i b'_j B(v_i, v'_j).$$

This formula shows that the pairings $B(v_i, v'_j) \in W$ do uniquely determine B. Conversely, given any $w_{ij} \in W$ define the set-theoretic map $B : V \times V' \to W$ by the condition $B(v, v') = \sum a_i b'_j w_{ij}$ where $v = \sum a_i v_i$ and $v' = \sum b'_j v'_j$ are the unique expansions of v and v' with respect to our choices of ordered bases of V and V'. The existence and uniqueness of such expansions for v and v' ensure that B is well-defined as a set-theoretic map, and obviously $B(v_i, v'_j) = w_{ij}$. It remains to check that this B we have just defined is bilienar, and this will be a computation with its definition. We check linearity in v' for a fixed $v \in V$, and the other way around goes similarly: for $v', \widetilde{v}' \in V'$ and $c', \widetilde{c}' \in F$ with $v' = \sum b'_j v'_j$ and $\widetilde{v}' = \sum \widetilde{b}'_j v'_j$ we have $c'v' + \widetilde{c}'\widetilde{v}' = \sum (c'b'_j + \widetilde{c}'\widetilde{b}'_j)v'_j$, and so

$$B(v, c'v' + \widetilde{c}'\widetilde{v}') = \sum_{i,j} a_i(c'b'_j + \widetilde{c}'\widetilde{b}'_j)w_{ij} = c' \cdot \sum_{i,j} a_i b'_j w_{ij} + \widetilde{c}' \cdot \sum_{i,j} a_i \widetilde{b}'_j w_{ij} = c'B(v, v') + \widetilde{c}B(v, \widetilde{v}'),$$

as desired. ∎

This lemma shows that the data of a bilinear pairing $B : V \times V' \to W$ is "the same" as the data of nn' vectors $w_{ij} \in W$ with $1 \leq i \leq n$ and $1 \leq j \leq n'$. But note: although the set of bilinear pairings $B : V \times V' \to W$ has nothing to do with the choices of bases, and the set of indexed choices of nn' vectors $w_{ij} \in W$ has nothing to do with the choices of bases, the bijection between these sets *depends on the bases*! This is actually not a surprise: it is analogous to the well-known fact from linear algebra that $\mathrm{Hom}(V, V')$ can be identified with the set $\mathrm{Mat}_{n' \times n}(F)$ of $n' \times n$ matrices over F upon choosing ordered bases of V and V', and that if we change the choices then the correspondence changes (even though the sets being put in correspondence, $\mathrm{Hom}(V, V')$ and $\mathrm{Mat}_{n' \times n}(F)$, do not "know" about choices of bases of V or V').

As an example, if we replace v_i with $-v_i$ for all i but we leave the v'_j's unchanged then under the modified correspondence resulting from this new ordered basis of V the new B associated to a given collection of w_{ij}'s is the *negative* of the old B (essentially because $\beta(-v_i, v'_j) = -\beta(v_i, v'_j)$ for any bilinear $\beta : V \times V' \to W$). Of course, if one makes a more elaborate change of bases then it becomes more complicated to describe how the correspondence between B's and $\{w_{ij}\}$'s changes (just as it can be messy to describe how the matrix corresponding to a linear map changes when we make a complicated change of bases on the source and target of the bilinear map). As but one more example, if $n \geq 2$ and we replace v_1 with $v_1 + v_2$ then since $B(v_1 + v_2, v'_j) = B(v_1, v'_j) + B(v_2, v'_j)$ we see that if B corresponds to $\{w_{ij}\}$ with respect to the first choice of ordered bases of V and V' then with respect to this modified basis for V the pairing B will correspond to the collection of vectors in W whose ij-choice remains w_{ij} for $i > 1$ but whose $1j$-vector is $w_{1j} + w_{2j}$ for each j. The upshot is that the bijection between the B's and the $\{w_{ij}\}$'s is *not* intrinsic to the triple of vector spaces V, V', and W, but rather depends on the choices of ordered bases of V and V'.

Before we apply the preceding considerations to build tensor products, it is convenient to slightly generalize the construction of the usual Euclidean space F^n. For any finite non-empty set S, we define the *Euclidean space on S over F* to be the set F^S of F-valued functions $S \to F$, given pointwise F-vector space structure; elements of F^S are typically denoted $(c_s)_{s \in S}$, or simply (c_s), with $c_s \in F$. For example, if $S = \{1, \dots, n\}$ then F^S is just the old example F^n (with $(c_1, \dots, c_n)$ corresponding to the function $c : \{1, \dots, n\} \to F$ sending i to c_i). For $S = \{1, \dots, n\}$, elements of F^S may be "visualized" as rows or columns of elements in F, due to the fact that the set $\{1, \dots, n\}$ has an order structure, but this ordering is logically irrelevant to the underlying linear structure (and it only intervenes due to how human beings prefer to write things conveniently on a piece of paper). The vector space F^S has dimension equal to the positive size of the non-empty set S, and it has a *standard basis* consisting of those functions $\delta_{s_0} : S \to F$ that send some $s_0 \in S$ to 1 and all other $s \in S$ to 0. (Explicitly, $(c_s)_{s \in S} = \sum_{s \in S} c_s \delta_s$.) In the special case $S = \{1, \dots, n\}$ this recovers the usual notion of standard basis for F^n.

2. The construction

Theorem 2.1. *Let V and V' be finite-dimensional vector spaces over F. A tensor product of V and V' exists.*

Proof. If V or V' vanish then we have seen that the vanishing bilinear pairing into the zero vector space does the job, so we now suppose $n = \dim V$ and $n' = \dim V'$ are positive. Let

$$S = \{1, \dots, n\} \times \{1, \dots, n'\}$$

be the set of ordered pairs (i, j) of integers with $1 \le i \le n$ and $1 \le j \le n'$.

Now choose ordered bases $\{v_i\}$ of V and $\{v'_j\}$ of V'. Let $T = F^S$, and let $t : V \times V' \to T$ be the set-theoretic map

$$(v, v') \mapsto (a_i b'_j)_{(i,j)\in S} \in F^S$$

where $v = \sum a_i v_i$ and $v' = \sum b'_j v'_j$ are the unique basis expansions of v and v'. Roughly speaking, t is just a listing of the set of pairwise products of coefficients of v and v' with respect to the choices of bases. (By the proof of Lemma 1.1, such products of coefficients are "all" we need to know to compute $B(v, v')$ via a universal formula in terms of the $B(v_i, v'_j)$'s, so it should not be surprising that t "knows" every B, as we shall see.)

Note that t is bilinear. For example, if $\widetilde{v} = \sum \widetilde{a}_i v_i$ is an element of V then for any $c, \widetilde{c} \in F$ we have $cv + \widetilde{c}\widetilde{v} = \sum (ca_i + \widetilde{c}\widetilde{a}_i)v_i$, and hence

$$t(cv + \widetilde{c}\widetilde{v}, v') = ((ca_i + \widetilde{c}\widetilde{a}_i)b'_j)_{(i,j)} = c(a_i b'_j)_{(i,j)} + \widetilde{c}(\widetilde{a}_i b'_j)_{(i,j)} = ct(v, v') + \widetilde{c}t(\widetilde{v}, v').$$

This gives linearity in the first slot when the second is fixed, and the same method works the other way around. Clearly $t(v_{i_0}, v'_j j_0) = \delta_{(i_0,j_0)} \in F^S = T$ is the element whose (i_0, j_0)-coordinate is 1 and whose other coordinates are zero. That is, the vectors $t(v_i, v'_j) \in F^S$ form the standard basis of F^S.

Now we verify the universal mapping property. Let $B : V \times V' \to W$ be a bilinear mapping. We need to prove the existence *and uniqueness* of a linear map $L : T \to W$ such that $L \circ t = B$, which is to say $L(t(v, v')) = B(v, v')$ for all $v \in V$ and $v' \in V'$. But whatever L is to be, $L \circ t$ and B are two W-valued bilinear pairings between V and V', so to verify their equality it is necessary and sufficient for them to agree on pairs of vectors from bases of V and V'. That is, the condition on L is exactly $L(t(v_i, v'_j)) = B(v_i, v'_j)$ for all i and j. But the vectors $t(v_i, v'_j) = \delta_{(i,j)} \in F^S$ form the standard basis, and so the condition on the linear map $L : F^S \to W$ is precisely that its value on the standard basis vector $\delta_{(i,j)}$ is $B(v_i, v'_j) \in W$ for each i and j. We know that to give a linear map from a finite-dimensional vector space to another vector space, it is necessary and sufficient to specify the values of the mapping on elements of a fixed basis, and so L exists and is uniquely determined by the conditions $\delta_{(i,j)} \mapsto B(v_i, v'_j)$. ■

This proof deserves to be understood once, but it should never be referred to again because although the underlying vector space F^S in the proof has nothing to do with choices of bases, the specific universal bilinear pairing $t : V \times V' \to F^S$ that is built in the proof depends very much on choices of bases. One should really imagine the tensor product $(V \otimes V', \otimes)$ as an abstract notion, and the proof merely provides a way to make a concrete model of it (or rather, a linear isomorphism $V \otimes V' \simeq F^S$) by using bases, much as the "abstract" dual space $V^\vee$ takes on the more concrete appearance of "row vectors" upon using a dual basis for a choice of basis of V. In the context of the above construction, different choices of bases simply give rise to different linear isomorphisms $V \otimes V' \simeq F^S$ carrying the universal pairing $V \times V' \to V \otimes V'$ to some specific bilinear pairing $V \times V' \to F^S$. Another example to keep in mind is the distinction between the abstract space $\mathrm{Hom}(V, V')$ and the concrete vector space of $n' \times n$ matrices: if we change bases, the identification with the concrete-looking space of matrices will change. A big difference between the notions of $\mathrm{Hom}(V, V')$ and $V^\vee$ and the notion of $V \otimes V'$ is that the first two can be constructed without mentioning bases whereas for $V \otimes V'$ we have only the construction with bases and thus it is the universal mapping property that provides us with the only basis-free way to think about (and work with) this space. (As we have noted above, there *is* actually a method to construct $V \otimes V'$ without making any choices at all, but the construction is so horrible and unworkable that it is uninformative when using tensor product spaces and trying to understand their properties.)

3. An example

Let us work out one concrete example of how the basis change impacts the explicit construction of the tensor product, and see how the data of the universal bilinear pairing actually dictates a unique linear isomorphism between the resulting constructions of the tensor product.

Let $V = F^2$ and $V' = F^3$, and let $\{e_1, e_2\}$ and $\{e'_1, e'_2, e'_3\}$ be the standard bases. Let $S = \{1, 2\} \times \{1, 2, 3\}$. Following the proof, $V \otimes V' = F^S$ with universal bilinear pairing $t : V \times V' \to F^S$ determined (in the sense of Lemma 1.1) by the conditions $t(e_i, e'_j) = \delta_{(i,j)}$, where $\delta_{(i,j)} \in F^S$ has ij-coordinate 1 and all other coordinates 0.

Now suppose we pick different ordered bases, such as

$$\{v_1 = e_1 + e_2, v_2 = -e_2\},\ \ \{v'_1 = e'_1, v'_2 = e'_2 - e'_3, v'_3 = 4e'_1 + e'_2\}.$$

Using the proof with these bases once again exhibits F^S as a tensor product, but now the bilinear pairing $\tau : V \times V' \to F^S$ that comes out of the proof is determined by the conditions $\tau(v_i, v'_j) = \delta_{(i,j)}$. Note that as bilinear pairings $V \times V' \to F^S$, t and τ are *not the same*! The uniqueness of tensor products up to unique isomorphism ensures that there must be a unique isomorphism between the pairs (F^S, t) and (F^S, τ), which is to say that there exists a unique F-linear isomorphism $L : F^S \simeq F^S$ such that $\tau = L \circ t$. The problem we want to solve is: what is L?

The equality $\tau = L \circ t$ for bilinear pairings $V \times V' \to F^S$ may be checked on pairs of basis vectors (by Lemma 1.1), and so we may use the bases $\{v_i\}$ and $\{v'_j\}$ as in the construction of τ. Thus, L has to satisfy

$$\delta_{(i,j)} = \tau(v_i, v'_j) = L(t(v_i, v'_j))$$

for all i, j. Using the definitions of the v_i's and v'_j's in terms of the bases $\{e_1, e_2\}$ and $\{e'_1, e'_2, e'_3\}$ of V and V' that were used in the construction of t, we have

$$t(v_1, v'_1) = t(e_1 + e_2, e'_1) = t(e_1, e'_1) + t(e_2, e'_1) = \delta_{(1,1)} + \delta_{(2,1)},$$

$$t(v_1, v'_2) = t(e_1 + e_2, e'_2 - e'_3) = t(e_1, e'_2) - t(e_1, e'_3) + t(e_2, e'_2) - t(e_2, e'_3) = \delta_{(1,2)} - \delta_{(1,3)} + \delta_{(2,2)} - \delta_{(2,3)},$$

and

$$t(v_1, v'_3) = t(e_1 + e_2, 4e'_1 + e'_2) = 4\delta_{(1,1)} + \delta_{(1,2)} + 4\delta_{(2,1)} + \delta_{(2,2)}.$$

This computes $t(v_1, v'_j)$ for all j, and similarly one computes $t(v_2, v'_j)$ for all j:

$$t(v_2, v'_1) = -\delta_{(2,1)},\ t(v_2, v'_2) = -\delta_{(2,2)} + \delta_{(2,3)},\ t(v_2, v'_3) = \delta_{(2,2)} + 4\delta_{(2,3)}.$$

We have now expressed the $t(v_i, v'_j)$'s in terms of the "standard basis" $\delta_{(1,1)}, \delta_{(1,2)}, \ldots$ of F^S, and so the conditions $\delta_{(i,j)} = L(t(v_i, v'_j))$ for all i and j become a system of linear equations on the matrix of $L : F^S \simeq F^S$ with respect to this basis of F^S. More specifically, in order to work with matrices we need to *order* the set S, and we choose the lexicographical ordering ($(i, j) > (i', j')$ if either $i > i'$ or $i = i'$ and $j > j'$). In this way, one see that the matrix of coefficients for the

$t(v_i, v'_j)$'s in terms of the $\delta_{(i,j)}$'s is

$$M = \begin{pmatrix} 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & -1 & 0 & 1 & -1 \\ 4 & 1 & 0 & 4 & 1 & 0 \\ 0 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 4 \end{pmatrix}$$

The conditions $\delta_{(i,j)} = L(t(v_i, v'_j))$ therefore become: the matrix of L with respect to the standard basis of F^S using the lexicographical ordering is the inverse of the matrix M. That computes L! This illustrates an important principle when working with tensor products: when using bases, sometimes things can get quite messy. It is fortunate, as we shall see, that one can virtually always work with tensor products without ever using a basis.

To summarize: using the matrix inverse to M and the standard basis of F^S with lexicographical ordering so as to convert 6×6 matrices into linear self-maps of F^S, we get a linear isomorphism $L : F^S \simeq F^S$ that carries t to τ and thereby explains the relationship between the two concrete models (F^S, t) and (F^S, τ) for the abstract tensor product $(V \otimes V', \otimes)$. Of course, the matrix M defines the inverse isomorphism $F^S \simeq F^S$ that carries τ to t.