Basic Algebra

Final Vresion, August, 2006
For Publication by Birkhäuser Boston
Along with a Companion Volume *Advanced Algebra*
In the Series

Cornerstones

Selected Pages from Chapter II: pp. 33, 50–65

Anthony W. Knapp

CHAPTER II

Vector Spaces over $\mathbb{Q}$, $\mathbb{R}$, and $\mathbb{C}$

Abstract. This chapter introduces vector spaces and linear maps between them, and it goes on to develop certain constructions of new vector spaces out of old, as well as various properties of determinants.

Sections 1–2 define vector spaces, spanning, linear independence, bases, and dimension. The sections make use of row reduction to establish dimension formulas for certain vector spaces associated with matrices. They conclude by stressing methods of calculation that have quietly been developed in proofs.

Section 3 relates matrices and linear maps to each other, first in the case that the linear map carries column vectors to column vectors and then in the general finite-dimensional case. Techniques are developed for working with the matrix of a linear map relative to specified bases and for changing bases. The section concludes with a discussion of isomorphisms of vector spaces.

Sections 4–6 take up constructions of new vector spaces out of old ones, together with corresponding constructions for linear maps. The four constructions of vector spaces in these sections are those of the dual of a vector space, the quotient of two vector spaces, and the direct sum and direct product of two or more vector spaces.

Section 7 introduces determinants of square matrices, together with their calculation and properties. Some of the results that are established are expansion in cofactors, Cramer's rule, and the value of the determinant of a Vandermonde matrix. It is shown that the determinant function is well defined on any linear map from a finite-dimensional vector space to itself.

Section 8 introduces eigenvectors and eigenvalues for matrices, along with their computation. Also, in this section the characteristic polynomial and the trace of a square matrix are defined, and all these notions are reinterpreted in terms of linear maps.

Section 9 proves the existence of bases for infinite-dimensional vector spaces and discusses the extent to which the material of the first eight sections extends from the finite-dimensional case to be valid in the infinite-dimensional case.

1. Spanning, Linear Independence, and Bases

This chapter develops a theory of rational, real, and complex vector spaces. Many readers will already be familiar with some aspects of this theory, particularly in the case of the vector spaces $\mathbb{Q}^n$, $\mathbb{R}^n$, and $\mathbb{C}^n$ of column vectors, where the tools developed from row reduction allow one to introduce geometric notions and to view geometrically the set of solutions to a set of linear equations. Thus we shall

Pages 34–49 do not appear in this file.

The proposition does not mean that one should necessarily be eager to make the identification of two vector spaces that are isomorphic. An important distinction is the one between "isomorphic" and "isomorphic via a canonically constructed linear map." The isomorphism of linear maps with matrices given by $L \mapsto \begin{pmatrix} L \\ \Delta\Gamma \end{pmatrix}$ is canonical since no choices are involved once Γ and Δ have been specified. This is a useful isomorphism because we can track matters down and use the isomorphism to make computations. On the other hand, it is not very useful to say merely that $\mathrm{Hom}_{\mathbb{F}}(U, V)$ and $M_{kn}(\mathbb{F})$ are isomorphic because they have the same dimension.

What tends to happen in practice is that vector spaces in applications come equipped with additional structure—some rigid geometry, or a multiplication operation, or something else. A general vector-space isomorphism has little chance of having any connection to the additional structure and thereby of being very helpful. On the other hand, a concrete isomorphism that is built by taking this additional structure into account may indeed be useful.

In the next section we shall encounter an example of an additional structure that involves neither a rigid geometry nor a multiplication operation. We shall introduce the "dual" V' of a vector space V, and we shall see that V and V' have the same dimension if V is finite-dimensional. But no particular isomorphism of V with V' is singled out as better than other ones, and it is wise not to try to identify these spaces. By contrast, the double dual V'' of V, which too will be constructed in the next section, will be seen to be isomorphic to V in the finite-dimensional case via a linear map $\iota : V \to V''$ that we define explicitly. The function ι is an example of a canonical isomorphism that we might want to exploit.

4. Dual Spaces

Let V be a vector space over $\mathbb{F}$. A **linear functional** on V is a linear map from V into $\mathbb{F}$. The space of all such linear maps, as we saw in Section 3, is a vector space. We denote it by V' and call it the **dual space** of V.

The development of Section 3 tells us right away how to compute the dual space of the space of column vectors $\mathbb{F}^n$. If Σ is the standard ordered basis of $\mathbb{F}^n$ and if 1 denotes the basis of $\mathbb{F}$ consisting of the scalar 1, then we can associate to a linear functional v' on $\mathbb{F}^n$ its matrix

$$\begin{pmatrix} v' \\ 1\Sigma \end{pmatrix} = (\, v'(e_1) \quad v'(e_2) \quad \cdots \quad v'(e_n)\,),$$

which is an n-dimensional row vector. The operation of v' on a column vector

$v = \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix}$ is given by Theorem 2.14. Namely, $v'(v)$ is a multiple of the scalar 1, and the theorem tells us how to compute this multiple:

$$\begin{pmatrix} v'(v) \\ 1 \end{pmatrix} = \begin{pmatrix} v' \\ 1\Sigma \end{pmatrix} \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} = \begin{pmatrix} v'(e_1) & v'(e_2) & \cdots & v'(e_n) \end{pmatrix} \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix}.$$

Thus the space of all linear functionals on $\mathbb{F}^n$ may be identified with the space of all n-dimensional row vectors, and the effect of the row vector on a column vector is given by matrix multiplication. Since the standard ordered basis of $\mathbb{F}^n$ and the basis 1 of $\mathbb{F}$ are singled out as special, this identification is actually canonical, and it is thus customary to make this identification without further comment.

For a more general vector space V, no natural way of writing down elements of V' comes to mind. Indeed, if a concrete V is given, it can help considerably in understanding V to have an identification of V' that does not involve choices. For example, in real analysis one proves in a suitable infinite-dimensional setting that a (continuous) linear functional on the space of integrable functions is given by integration with a bounded function, and that fact simplifies the handling of the space of integrable functions.

In any event, the canonical identification of linear functionals that we found for $\mathbb{F}^n$ does not work once we pass to a more general finite-dimensional vector space V. To make such an identification in the absence of additional structure, we first fix an ordered basis $(v_1, \ldots, v_n)$ of V. If we do so, then V' is indeed identified with the space of n-dimensional row vectors. The members of V' that correspond to the standard basis of row vectors, i.e., the row vectors that are 1 in one entry and are 0 elsewhere, are of special interest. These are the linear functionals v_i' such that

$$v_i'(v_j) = \delta_{ij},$$

where δ_{ij} is the Kronecker delta. Since these standard row vectors form a basis of the space of row vectors, $(v_1', \ldots, v_n')$ is an ordered basis of V'. If the members of the ordered basis $(v_1, \ldots, v_n)$ are permuted in some way, the members of $(v_1', \ldots, v_n')$ are permuted in the same way. Thus the basis $\{v_1', \ldots, v_n'\}$ depends only on the basis $\{v_1, \ldots, v_n\}$, not on the enumeration.[6] The basis $\{v_1', \ldots, v_n'\}$ is called the **dual basis** of V relative to $\{v_1, \ldots, v_n\}$. A consequence of this discussion is the following result.

Proposition 2.19. If V is a finite-dimensional vector space with dual V', then V' is finite-dimensional with $\dim V' = \dim V$.

[6] Although the enumeration is not important, more structure is present here than simply an association of an unordered basis of V' to an unordered basis of V. Each member of $\{v_1', \ldots, v_n'\}$ is matched to a particular member of $\{v_1, \ldots, v_n\}$, namely the one on which it takes the value 1.

Linear functionals play an important role in working with a vector space. To understand this role, it is helpful to think somewhat geometrically. Imagine the problem of describing a vector subspace of a given vector space. One way of describing it is from the inside, so to speak, by giving a spanning set. In this case we end up by describing the subspace in terms of parameters, the parameters being the scalar coefficients when we say that the subspace is the set of all finite linear combinations of members of the spanning set. Another way of describing the subspace is from the outside, cutting it down by conditions imposed on its elements. These conditions tend to be linear equations, saying that certain linear maps on the elements of the subspace give 0. Typically the subspace is then described as the intersection of the kernels of some set of linear maps. Frequently these linear maps will be scalar-valued, and then we are in a situation of describing the subspace by a set of linear functionals.

We know that every vector subspace of a finite-dimensional vector space V can be described from the inside in this way; we merely give all its members. A statement with more content is that we can describe it with finitely many members; we can do so because we know that every vector subspace of V has a basis.

For linear functionals really to be useful, we would like to know a corresponding fact about describing subspaces from the outside—that every vector subspace U of a finite-dimensional V can be described as the intersection of the kernels of a finite set of linear functionals. To do so is easy. We take a basis of the vector subspace U, say $\{v_1, \dots, v_r\}$, extend it to a basis of V by adjoining vectors $v_{r+1}, \dots, v_n$, and form the dual basis $\{v_1', \dots, v_n'\}$ of V'. The subspace U is then described as the set of all vectors v in V such that $v_j'(v) = 0$ for $r+1 \le j \le n$. The following proposition expresses this fact in ways that are independent of the choice of a basis. It uses the terminology **annihilator** of U, denoted by $\mathrm{Ann}(U)$, for the vector subspace of all members v' of V' with $v'(u) = 0$ for all u in U.

Proposition 2.20. Let V be a finite-dimensional vector space, and let U be a vector subspace of V. Then

(a) $\dim U + \dim \mathrm{Ann}(U) = \dim V$,
(b) every linear functional on U extends to a linear functional on V,
(c) whenever v_0 is a member of V that is not in U, there exists a linear functional on V that is 0 on U and is 1 on v_0.

PROOF. We retain the notation above, writing $\{v_1, \dots, v_r\}$ for a basis of U, $v_{r+1}, \dots, v_n$ for vectors that are adjoined to form a basis of V, and $\{v_1', \dots, v_n'\}$ for the dual basis of V'. For (a), we check that $\{v_{r+1}', \dots, v_n'\}$ is a basis of $\mathrm{Ann}(U)$. It is enough to see that they span $\mathrm{Ann}(U)$. These linear functionals are 0 on every member of the basis $\{v_1, \dots, v_r\}$ of U and hence are in $\mathrm{Ann}(U)$. On the other hand, if v' is a member of $\mathrm{Ann}(U)$, we can certainly write $v' = c_1 v_1' + \cdots + c_n v_n'$

for some scalars $c_1, \ldots, c_n$. Since v' is 0 on U, we must have $v'(v_i) = 0$ for $i \le r$. Since $v'(v_i) = c_i$, we obtain $c_i = 0$ for $i \le r$. Therefore v' is a linear combination of $v'_{r+1}, \ldots, v'_n$, and (a) is proved.

For (b), let us observe that the restrictions $v'_1\big|_U, \ldots, v'_r\big|_U$ form the dual basis of U' relative to the basis $\{v_1, \ldots, v_r\}$ of U. If u' is in U', we can therefore write $u' = c_1 v'_1\big|_U + \cdots + c_r v'_r\big|_U$ for some scalars $c_1, \ldots, c_r$. Then $v' = c_1 v'_1 + \cdots + c_r v'_r$ is the required extension of u' to all of V.

For (c), we use a special choice of basis of V in the argument above. Namely, we still take $\{v_1, \ldots, v_r\}$ to be a basis of U, and then we let $v_{r+1} = v_0$. Finally we adjoin $v_{r+2}, \ldots, v_n$ to obtain a basis $\{v_1, \ldots, v_n\}$ of V. Then v'_{r+1} has the required property. $\square$

If $L : U \to V$ is a linear map between finite-dimensional vector spaces, then the formula

$$(L^t(v'))(u) = v'(L(u)) \qquad \text{for } u \in U \text{ and } v' \in V'$$

defines a linear map $L^t : V' \to U'$. The linear map L^t is called the **contragredient** of L. The matrix of the contragredient of L is the transpose of the matrix of L in the following sense.[7]

Proposition 2.21. Let $L : U \to V$ be a linear map between finite-dimensional vector spaces, let $L^t : V' \to U'$ be its contragredient, let Γ and Δ be respective ordered bases of U and V, and let Γ' and Δ' be their dual ordered bases. Then

$$\begin{pmatrix} L^t \\ \Gamma'\Delta' \end{pmatrix} = \begin{pmatrix} L \\ \Delta\Gamma \end{pmatrix}.$$

PROOF. Let $\Gamma = (u_1, \ldots, u_n)$, $\Delta = (v_1, \ldots, v_k)$, $\Gamma' = (u'_1, \ldots, u'_n)$, and $\Delta' = (v'_1, \ldots, v'_k)$. Write B and A for the respective matrices in the formula in question. The equations $L(u_j) = \sum_{i'=1}^k A_{i'j} v_{i'}$ and $L^t(v'_i) = \sum_{j'=1}^n B_{j'i} u'_{j'}$ imply that

$$v'_i(L(u_j)) = v'_i\big(\textstyle\sum_{i'=1}^k A_{i'j} v_{i'}\big) = A_{ij}$$

and

$$L^t(v'_i)(u_j) = \textstyle\sum_{j'=1}^n B_{j'i} u'_{j'}(u_j) = B_{ji}.$$

Therefore $B_{ji} = L^t(v'_i)(u_j) = v'_i(L(u_j)) = A_{ij}$, as required. $\square$

[7] A general principle is involved in the definition of contragredient once we have a definition of dual vector space, and we shall see further examples of this principle in the next two sections and in later chapters: whenever a new systematic construction appears for the objects under study, it is well to look for a corresponding construction with the functions relating these new objects. In language to be introduced near the end of Chapter IV, the context for the construction will be a "category," and the principle says that it is well to see whether the construction is that of a "functor" on the category.

With V finite-dimensional, now consider $V'' = (V')'$, the double dual. In the case that $V = \mathbb{F}^n$, we saw that V' could be viewed as the space of row vectors, and it is reasonable to expect V'' to involve a second transpose and again be the space of column vectors. If so, then V gets identified with V''. In fact, this is true in all cases, and we argue as follows. If v is in V, we can define a member $\iota(v)$ of V'' by

$$\iota(v)(v') = v'(v) \qquad \text{for } v \in V \text{ and } v' \in V'.$$

This definition makes sense whether or not V is finite-dimensional. The function ι is a linear map from V into V'' called the **canonical map** of V into V''. It is independent of any choice of basis.

Proposition 2.22. If V is any finite-dimensional vector space over $\mathbb{F}$, then the canonical map $\iota : V \to V''$ is one-one onto.

REMARKS. In the infinite-dimensional case the canonical map is one-one but it is not onto. The proof that it is one-one uses the fact that V has a basis, but we have deferred the proof of this fact about infinite-dimensional vector spaces to Section 9. Problem 14 at the end of the chapter will give an example of an infinite-dimensional V for which ι does not carry V onto V''. When combined with the first corollary in Section A6 of the appendix, this example shows that ι *never* carries V onto V'' in the infinite-dimensional case.

PROOF. We saw in Section 3 that a linear map ι is one-one if and only if $\ker \iota = 0$. Thus suppose $\iota(v) = 0$. Then $0 = \iota(v)(v') = v'(v)$ for all v'. Arguing by contradiction, suppose $v \neq 0$. Then we can extend $\{v\}$ to a basis of V, and the linear functional v' that is 1 on v and is 0 on the other members of the basis will have $v'(v) \neq 0$, contradiction. We conclude that ι is one-one. By Proposition 2.19 we have

$$\dim V = \dim V' = \dim V''. \tag{$*$}$$

Since ι is one-one, it carries any basis of V to a linearly independent set in V''. This linearly independent set has to be a basis, by Corollary 2.4 and the dimension formula $(*)$. □

5. Quotients of Vector Spaces

This section constructs a vector space V/U out of a vector space V and a vector subspace U. We begin with the example illustrated in Figure 2.1. In the vector space $V = \mathbb{R}^2$, let U be a line through the origin. The lines parallel to U are of the form $v + U = \{v + u \mid u \in U\}$, and we make the set of these lines into a vector space by defining $(v_1 + U) + (v_2 + U) = (v_1 + v_2) + U$ and

$c(v + U) = cv + U$. The figure suggests that if we were to take any other line W through the origin, then W would meet all the lines $v + U$, and the notion of addition of lines $v + U$ would correspond exactly to addition in W. Indeed we can successfully make such a correspondence, but the advantage of introducing the vector space of all lines $v + U$ is that it is canonical, independent of the kind of choice we have to make in selecting W. One example of the utility of having a canonical construction is the ease with which we obtain correspondence of linear maps stated in Proposition 2.25 below. Other examples will appear later.

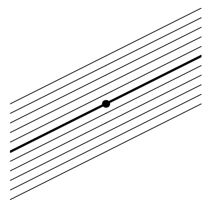

FIGURE 2.1. The vector space of lines $v + U$ in $\mathbb{R}^2$ parallel to a given line U through the origin.

Proposition 2.23. Let V be a vector space over $\mathbb{F}$, and let U be a vector subspace. The relation defined by saying that $v_1 \sim v_2$ if $v_1 - v_2$ is in U is an equivalence relation, and the equivalence classes are all sets of the form $v + U$ with $v \in V$. The set of equivalence classes V/U is a vector space under the definitions

$$(v_1 + U) + (v_2 + U) = (v_1 + v_2) + U,$$
$$c(v + U) = cv + U,$$

and the function $q(v) = v + U$ is linear from V onto V/U with kernel U.

REMARKS. We say that V/U is the **quotient space** of V by U. The linear map $q(v) = v + U$ is called the **quotient map** of V onto V/U.

PROOF. The properties of an equivalence relation are established as follows:

$v_1 \sim v_1$	because 0 is in U,
$v_1 \sim v_2$ implies $v_2 \sim v_1$	because U is closed under negatives,
$v_1 \sim v_2$ and $v_2 \sim v_3$ together imply $v_1 \sim v_3$	because U is closed under addition.

Thus we have equivalence classes. The class of v_1 consists of all vectors v_2 such that $v_2 - v_1$ is in U, hence consists of all vectors in $v_1 + U$. Thus the equivalence classes are indeed the sets $v + U$.

Let us check that addition and scalar multiplication, as given in the statement of the proposition, are well defined. For addition let $v_1 \sim w_1$ and $v_2 \sim w_2$. Then $v_1 - w_1$ and $v_2 - w_2$ are in U. Since U is a vector subspace, the sum $(v_1 - w_1) + (v_2 - w_2) = (v_1 + v_2) - (w_1 + w_2)$ is in U. Thus $v_1 + v_2 \sim w_1 + w_2$, and addition is well defined. For scalar multiplication let $v \sim w$, and let a scalar c be given. Then $v - w$ is in U, and $c(v - w) = cv - cw$ is in U since U is a vector subspace. Hence $cv \sim cw$, and scalar multiplication is well defined.

The vector-space properties of V/U are consequences of the properties for V. To illustrate, consider associativity of addition. The argument in this case is that

$$\begin{aligned}((v_1 + U) + (v_2 + U)) + (v_3 + U) &= ((v_1 + v_2) + U) + (v_3 + U) \\ &= ((v_1 + v_2) + v_3) + U = (v_1 + (v_2 + v_3)) + U \\ &= (v_1 + U) + ((v_2 + v_3) + U) = (v_1 + U) + ((v_2 + U) + (v_3 + U)).\end{aligned}$$

Finally the quotient map $q : V \to V/U$ given by $q(v) = v + U$ is certainly linear. Its kernel is $\{v \mid v + U = 0 + U\}$, and this equals $\{v \mid v \in U\}$, as asserted. The map q is onto V/U since $v + U = q(v)$. $\square$

Corollary 2.24. If V is a vector space over $\mathbb{F}$ and U is a vector subspace, then

(a) $\dim V = \dim U + \dim(V/U)$,

(b) the subspace U is the kernel of some linear map defined on V.

REMARK. The first conclusion is valid even when all the spaces are not finite-dimensional. For current purposes it is sufficient to regard $\dim V$ as $+\infty$ if V is infinite-dimensional; the sum of $+\infty$ and any dimension as $+\infty$.

PROOF. Let q be the quotient map. The linear map q meets the conditions of (b). For (a), take a basis of U and extend to a basis of V. Then the images under q of the additional vectors form a basis of V/U. $\square$

Quotients of vector spaces allow for the factorization of certain linear maps, as indicated in Proposition 2.25 and Figure 2.2.

Proposition 2.25. Let $L : V \to W$ be a linear map between vector spaces over $\mathbb{F}$, let $U_0 = \ker L$, let U be a vector subspace of V contained in U_0, and let $q : V \to V/U$ be the quotient map. Then there exists a linear map $\overline{L} : V/U \to W$ such that $L = \overline{L}q$. It has the same image as L, and $\ker \overline{L} = \{u_0 + U \mid u_0 \in U_0\}$.

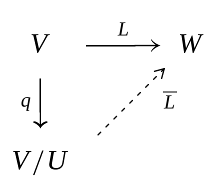

FIGURE 2.2. Factorization of linear maps via a quotient of vector spaces.

REMARK. One says that L **factors through** V/U or **descends** to V/U.

PROOF. The definition of $\overline{L}$ has to be $\overline{L}(v+U) = L(v)$. This forces $\overline{L}q = L$, and $\overline{L}$ will have to be linear. What needs proof is that $\overline{L}$ is well defined. Thus suppose $v_1 \sim v_2$. We are to prove that $\overline{L}(v_1 + U) = \overline{L}(v_2 + U)$, i.e., that $L(v_1) = L(v_2)$. Now $v_1 - v_2$ is in $U \subseteq U_0$, and hence $L(v_1 - v_2) = 0$. Then $L(v_1) = L(v_1 - v_2) + L(v_2) = L(v_2)$, as required. This proves that $\overline{L}$ is well defined, and the conclusions about the image and the kernel of $\overline{L}$ are immediate from the definition. □

Corollary 2.26. Let $L : V \to W$ be a linear map between vector spaces over $\mathbb{F}$, and suppose that L is onto W and has kernel U. Then V/U is canonically isomorphic to W.

PROOF. Take $U = U_0$ in Proposition 2.25, and form $\overline{L} : V/U \to W$ with $L = \overline{L}q$. The proposition shows that $\overline{L}$ is onto W and has trivial kernel, i.e., the 0 element of V/U. Having trivial kernel, $\overline{L}$ is one-one. □

Theorem 2.27 (First Isomorphism Theorem). Let $L : V \to W$ be a linear map between vector spaces over $\mathbb{F}$, and suppose that L is onto W and has kernel U. Then the map $S \mapsto L(S)$ gives a one-one correspondence between

(a) the vector subspaces S of V containing U and
(b) the vector subspaces of W.

REMARK. As in Section A1 of the appendix, we write $L(S)$ and $L^{-1}(T)$ to indicate the direct and inverse images of S and T, respectively.

PROOF. The passage from (a) to (b) is by direct image under L, and the passage from (b) to (a) will be by inverse image under L^{-1}. Certainly the direct image of a vector subspace as in (a) is a vector subspace as in (b). We are to show that the inverse image of a vector subspace as in (b) is a vector subspace as in (a) and that these two procedures invert one another.

For any vector subspace T of W, $L^{-1}(T)$ is a vector subspace of V. In fact, if v_1 and v_2 are in $L^{-1}(T)$, we can write $L(v_1) = t_1$ and $L(v_2) = t_2$ with t_1 and t_2 in T. Then the equations $L(v_1 + v_2) = t_1 + t_2$ and $L(cv_1) = cL(v_1) = ct_1$ show that $v_1 + v_2$ and cv_1 are in $L^{-1}(T)$.

Moreover, the vector subspace $L^{-1}(T)$ contains $L^{-1}(0) = U$. Therefore the inverse image under L of a vector subspace as in (b) is a vector subspace as in (a). Since L is a function, we have $L(L^{-1}(T)) = T$. Thus passing from (b) to (a) and back recovers the vector subspace of W.

If S is a vector subspace of V containing U, we still need to see that $S = L^{-1}(L(S))$. Certainly $S \subseteq L^{-1}(L(S))$. In the reverse direction let v be in $L^{-1}(L(S))$. Then $L(v)$ is in $L(S)$, i.e., $L(v) = L(s)$ for some s in S. Since L

is linear, $L(v-s)=0$. Thus $v-s$ is in $\ker L = U$, which is contained in S by assumption. Then s and $v-s$ are in S, and hence v is in S. We conclude that $L^{-1}(L(S)) \subseteq S$, and thus passing from (a) to (b) and then back recovers the vector subspace of V containing U.

If V is a vector space and V_1 and V_2 are vector subspaces, then we write V_1+V_2 for the set V_1+V_2 of all sums v_1+v_2 with $v_1 \in V_1$ and $v_2 \in V_2$. This is again a vector subspace of V and is called the **sum** of V_1 and V_2. If we have vector subspaces $V_1, \dots, V_n$, we abbreviate $((\cdots(V_1+V_2)+V_3)+\cdots+V_n)$ as $V_1+\cdots+V_n$. □

Theorem 2.28 (Second Isomorphism Theorem). Let M and N be vector subspaces of a vector space V over $\mathbb{F}$. Then the map $n+(M\cap N)\mapsto n+M$ is a well-defined canonical vector-space isomorphism

$$N/(M\cap N)\cong(M+N)/M.$$

PROOF. The function $L(n+(M\cap N))=n+M$ is well defined since $M\cap N\subseteq M$, and L is linear. The domain of L is $\{n+(M\cap N)\mid n\in N\}$, and the kernel is the subset of this where n lies in M as well as N. For this to happen, n must be in $M\cap N$, and thus the kernel is the 0 element of $N/(M\cap N)$. Hence L is one-one.

To see that L is onto $(M+N)/M$, let $(m+n)+M$ be given. Then $n+(M\cap N)$ maps to $n+M$, which equals $(m+n)+M$. Hence L is onto. □

Corollary 2.29. Let M and N be finite-dimensional vector subspaces of a vector space V over $\mathbb{F}$. Then

$$\dim(M+N)+\dim(M\cap N)=\dim M+\dim N.$$

PROOF. Theorem 2.28 and two applications of Corollary 2.24a yield

$$\begin{aligned}\dim(M+N)-\dim M &= \dim((M+N)/M)\\ &= \dim(N/(M\cap N)) = \dim N-\dim(M\cap N),\end{aligned}$$

and the result follows. □

6. Direct Sums and Direct Products of Vector Spaces

In this section we introduce the direct sum and direct product of two or more vector spaces over $\mathbb{F}$. When there are only finitely many such subspaces, these constructions come to the same thing, and we call it "direct sum." We begin with the case that two vector spaces are given.

We define two kinds of direct sums. The **external direct sum** of two vector spaces V_1 and V_2 over $\mathbb{F}$, written $V_1 \oplus V_2$, is a vector space obtained as follows. The underlying set is the set-theoretic product, i.e., the set $V_1 \times V_2$ of ordered pairs (v_1, v_2) with $v_1 \in V_1$ and $v_2 \in V_2$. The operations of addition and scalar multiplication are defined coordinate by coordinate:

$$(u_1, u_2) + (v_1, v_2) = (u_1 + v_1, u_2 + v_2),$$
$$c(v_1, v_2) = (cv_1, cv_2),$$

and it is immediate that $V_1 \oplus V_2$ satisfies the defining properties of a vector space.

If $\{a_i\}$ is a basis of V_1 and $\{b_j\}$ is a basis of V_2, then it follows from the formula $(v_1, v_2) = (v_1, 0) + (0, v_2)$ that $\{(a_i, 0)\} \cup \{(0, b_j)\}$ is a basis of $V_1 \oplus V_2$. Consequently if V_1 and V_2 are finite-dimensional, then $V_1 \oplus V_2$ is finite-dimensional with

$$\dim(V_1 \oplus V_2) = \dim V_1 + \dim V_2.$$

Associated to the construction of the external direct sum of two vector spaces are four linear maps of interest:

two "projections,"	$p_1 : V_1 \oplus V_2 \to V_1$	with $p_1(v_1, v_2) = v_1$,
	$p_2 : V_1 \oplus V_2 \to V_2$	with $p_2(v_1, v_2) = v_2$,
two "injections,"	$i_1 : V_1 \to V_1 \oplus V_2$	with $i_1(v_1) = (v_1, 0)$,
	$i_2 : V_2 \to V_1 \oplus V_2$	with $i_2(v_2) = (0, v_2)$.

These have the properties that

$$p_r i_s = \begin{cases} I & \text{on } V_s \text{ if } r = s, \\ 0 & \text{on } V_s \text{ if } r \neq s, \end{cases}$$

$$i_1 p_1 + i_2 p_2 = I \qquad \text{on } V_1 \oplus V_2.$$

The second notion of direct sum captures the idea of recognizing a situation as canonically isomorphic to an external direct sum. This is based on the following proposition.

Proposition 2.30. Let V be a vector space over $\mathbb{F}$, and let V_1 and V_2 be vector subspaces of V. Then the following conditions are equivalent:

(a) every member v of V decomposes uniquely as $v = v_1 + v_2$ with $v_1 \in V_1$ and $v_2 \in V_2$,

(b) $V_1 + V_2 = V$ and $V_1 \cap V_2 = 0$,

(c) the function from the external direct sum $V_1 \oplus V_2$ to V given by $(v_1, v_2) \mapsto v_1 + v_2$ is an isomorphism of vector spaces.

REMARKS.

(1) If V is a vector space with vector subspaces V_1 and V_2 satisfying the equivalent conditions of Proposition 2.30, then we say that V is the **internal direct sum** of V_1 and V_2. It is customary to write $V = V_1 \oplus V_2$ in this case even though what we have is a canonical isomorphism of the two sides, not an equality.

(2) The dimension formula

$$\dim(V_1 \oplus V_2) = \dim V_1 + \dim V_2$$

for an internal direct sum follows, on the one hand, from the corresponding formula for external direct sums; it follows, on the other hand, by using (b) and Corollary 2.29.

(3) In the proposition it is possible to establish a fourth equivalent condition as follows: there exist linear maps $p_1 : V \to V$, $p_2 : V \to V$, $i_1 : \text{image } p_1 \to V$, and $i_2 : \text{image } p_2 \to V$ such that

- $p_r i_s p_s$ equals p_r if $r = s$ and equals 0 if $r \neq s$,
- $i_1 p_1 + i_2 p_2 = I$, and
- $V_1 = \text{image } i_1 p_1$ and $V_2 = \text{image } i_2 p_2$.

PROOF. If (a) holds, then the existence of the decomposition $v = v_1 + v_2$ shows that $V_1 + V_2 = V$. If v is in $V_1 \cap V_2$, then $0 = v + (-v)$ is a decomposition of the kind in (a), and the uniqueness forces $v = 0$. Therefore $V_1 \cap V_2 = 0$. This proves (b).

The function in (c) is certainly linear. If (b) holds and v is given in V, then the identity $V_1 + V_2 = V$ allows us to decompose v as $v = v_1 + v_2$. This proves that the linear map in (c) is onto. To see that it is one-one, suppose that $v_1 + v_2 = 0$. Then $v_1 = -v_2$ shows that v_1 is in $V_1 \cap V_2$. By (b), this intersection is 0. Therefore $v_1 = v_2 = 0$, and the linear map in (c) is one-one.

If (c) holds, then the fact that the linear map in (c) is onto V proves the existence of the decomposition in (a). For uniqueness, suppose that $v_1 + v_2 = u_1 + u_2$ with u_1 and v_1 in V_1 and with u_2 and v_2 in V_2. Then (u_1, u_2) and (v_1, v_2) have the same image under the linear map in (c). Since the function in (c) is assumed one-one, we conclude that $(u_1, u_2) = (v_1, v_2)$. This proves the uniqueness of the decomposition in (a). □

If $V = V_1 \oplus V_2$ is a direct sum, then we can use the above projections and injections to pass back and forth between linear maps with V_1 and V_2 as domain or range and linear maps with V as domain or range. This passage back and forth is called the **universal mapping property** of $V_1 \oplus V_2$ and will be seen later in this section to characterize $V_1 \oplus V_2$ up to canonical isomorphism. Let us be specific about how this property works.

To arrange for V to be the range, suppose that U is a vector space over $\mathbb{F}$ and that $L_1 : U \to V_1$ and $L_2 : U \to V_2$ are linear maps. Then we can define a linear map $L : U \to V$ by $L = i_1 L_1 + i_2 L_2$, i.e., by

$$L(u) = (i_1 L_1 + i_2 L_2)(u) = (L_1(u), L_2(u)),$$

and we can recover L_1 and L_2 from L by $L_1 = p_1 L$ and $L_2 = p_2 L$.

To arrange for V to be the domain, suppose that W is a vector space over $\mathbb{F}$ and that $M_1 : V_1 \to W$ and $M_2 : V_2 \to W$ are linear maps. Then we can define a linear map $M : V \to W$ by $M = M_1 p_1 + M_2 p_2$, i.e., by

$$M(v_1, v_2) = M_1(v_1) + M_2(v_2),$$

and we can recover M_1 and M_2 from M by $M_1 = M i_1$ and $M_2 = M i_2$.

The notion of direct sum readily extends to the direct sum of n vector spaces over $\mathbb{F}$. The **external direct sum** $V_1 \oplus \cdots \oplus V_n$ is the set of ordered pairs $(v_1, \dots, v_n)$ with each v_j in V_j and with addition and scalar multiplication defined coordinate by coordinate. In the finite-dimensional case we have

$$\dim(V_1 \oplus \cdots \oplus V_n) = \dim V_1 + \cdots + \dim V_n.$$

If $V_1, \dots, V_n$ are given as vector subspaces of a vector space V, then we say that V is the **internal direct sum** of $V_1, \dots, V_n$ if the equivalent conditions of Proposition 2.31 below are satisfied. In this case we write $V = V_1 \oplus \cdots \oplus V_n$ even though once again we really have a canonical isomorphism rather than an equality.

Proposition 2.31. Let V be a vector space over $\mathbb{F}$, and let $V_1, \dots, V_n$ be vector subspaces of V. Then the following conditions are equivalent:

(a) every member v of V decomposes uniquely as $v = v_1 + \cdots + v_n$ with $v_j \in V_j$ for $1 \le j \le n$,

(b) $V_1 + \cdots + V_n = V$ and also $V_j \cap (V_1 + \cdots + V_{j-1} + V_{j+1} + \cdots + V_n) = 0$ for each j with $1 \le j \le n$,

(c) the function from the external direct sum $V_1 \oplus \cdots \oplus V_n$ to V given by $(v_1, \dots, v_n) \mapsto v_1 + \cdots + v_n$ is an isomorphism of vector spaces.

Proposition 2.31 is proved in the same way as Proposition 2.30, and the expected analog of Remark 3 with that proposition is valid as well. Notice that the second condition in (b) is stronger than the condition that $V_i \cap V_j = 0$ for all $i \neq j$. Figure 2.3 illustrates how the condition $V_i \cap V_j = 0$ for all $i \neq j$ can be satisfied even though (b) is not satisfied and even though the vector subspaces do not therefore form a direct sum.

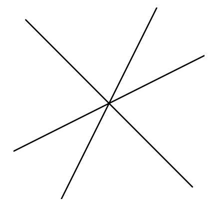

FIGURE 2.3. Three 1-dimensional vector subspaces of $\mathbb{R}^2$ such that each pair has intersection 0.

If $V = V_1 \oplus \cdots \oplus V_n$ is a direct sum, then we can define projections $p_1, \ldots, p_n$ and injections $i_1, \ldots, i_n$ in the expected way, and we again get a universal mapping property. That is, we can pass back and forth between linear maps with $V_1, \ldots, V_n$ as domain or range and linear maps with V as domain or range. The argument given above for $n = 2$ is easily adjusted to handle general n, and we omit the details.

To generalize the above notions to infinitely many vector spaces, there are two quite different ways of proceeding. Let us treat first the external constructions. Let a nonempty collection of vector spaces V_α over $\mathbb{F}$ be given, one for each $\alpha \in A$. The **direct sum** $\bigoplus_{\alpha\in A} V_\alpha$ is the set of all tuples $\{v_\alpha\}$ in the Cartesian product $\times_{\alpha\in A} V_\alpha$ with all but finitely many v_α equal to 0 and with addition and scalar multiplication defined coordinate by coordinate. For this construction we obtain a basis as the union of embedded bases of the constituent spaces. The **direct product** $\prod_{\alpha\in A} V_\alpha$ is the set of *all* tuples $\{v_\alpha\}$ in $\times_{\alpha\in A} V_\alpha$, again with addition and scalar multiplication defined coordinate by coordinate. When there are only finitely many factors $V_1, \ldots, V_n$, the direct product, which manifestly coincides with the direct sum, is sometimes denoted by $V_1 \times \cdots \times V_n$. For the direct product when there are infinitely many factors, there is no evident way to obtain a basis of the product from bases of the constituents.

The projections and injections that we defined in the case of finitely many vector spaces are still meaningful here. The universal mapping property is still valid as well, but it splinters into one form for direct sums and another form for direct products. The formulas given above for using linear maps with the V_α's as domain or range to define linear maps with the direct sum or direct product as domain or range may involve sums with infinitely many nonzero terms, and they are not directly usable. Instead, the formulas that continue to make sense are the ones for recovering linear maps with the V_α's as domain or range from linear maps with the direct sum or direct product as domain or range. These turn out to determine the formulas uniquely for the linear maps with the direct sum or direct product as domain or range. In other words, the appropriate universal mapping property uniquely determines the direct sum or direct product up to an isomorphism that respects the relevant projections and injections.

Let us see to the details. We denote typical members of $\prod_{\alpha\in A} V_\alpha$ and $\bigoplus_{\alpha\in A} V_\alpha$ by $\{v_\alpha\}_{\alpha\in A}$, with the understanding that only finitely many v_α can be nonzero in the case of the direct sum. The formulas are

$$p_\beta : \prod_{\alpha\in A} V_\alpha \to V_\beta \qquad \text{with } p_\beta\big(\{v_\alpha\}_{\alpha\in A}\big) = v_\beta,$$

$$i_\beta : V_\beta \to \bigoplus_{\alpha\in A} V_\alpha \qquad \text{with } i_\beta(v_\beta) = \{w_\alpha\}_{\alpha\in A} \text{ and } w_\alpha = \begin{cases} v_\beta & \text{if } \alpha = \beta, \\ 0 & \text{if } \alpha \neq \beta. \end{cases}$$

If U is a vector space over $\mathbb{F}$ and if a linear map $L_\beta : U \to V_\beta$ is given for each $\beta \in A$, we can obtain a linear map $L : U \to \prod_{\alpha\in A} V_\alpha$ that satisfies $p_\beta L = L_\beta$ for all β. The definition that makes perfectly good sense is

$$L(u) = \{L(u)_\alpha\}_{\alpha\in A} = \{L_\alpha(u)\}_{\alpha\in A}.$$

What does not make sense is to try to express the right side in terms of the injections i_α; we cannot write the right side as $\sum_{\alpha\in A} i_\alpha(L_\alpha(u))$ because infinitely many terms might be nonzero.

If W is a vector space and a linear map $M_\beta : V_\beta \to W$ is given for each β, we can obtain a linear map $M : \bigoplus_{\alpha\in A} V_\alpha \to W$ that satisfies $Mi_\beta = M_\beta$ for all β; the definition that makes perfectly good sense is

$$M\big(\{v_\alpha\}_{\alpha\in A}\big) = \sum_{\alpha\in A} M_\alpha(v_\alpha).$$

The right side is meaningful since only finitely many v_α can be nonzero. It can be misleading to write the formula as $M = \sum_{\alpha\in A} M_\alpha p_\alpha$ because infinitely many of the linear maps $M_\alpha p_\alpha$ can be nonzero functions.

In any event, we have a universal mapping property in both cases—for the direct product with the projections in place and for the direct sum with the injections in place. Let us see that these universal mapping properties characterize direct products and direct sums up to an isomorphism respecting the projections and injections, and that they allow us to define and recognize "internal" direct products and direct sums.

A **direct product** of a set of vector spaces V_α over $\mathbb{F}$ for $\alpha \in A$ consists of a vector space V and a system of linear maps $p_\alpha : V \to V_\alpha$ with the following **universal mapping property**: whenever U is a vector space and $\{L_\alpha\}$ is a system of linear maps $L_\alpha : U \to V_\alpha$, then there exists a unique linear map $L : U \to V$ such that $p_\alpha L = L_\alpha$ for all α. See Figure 2.4. The external direct product establishes existence of a direct product, and Proposition 2.32 below establishes its uniqueness up to an isomorphism of the V's that respects the p_α's. A direct product is said to be **internal** if each V_α is a vector subspace of V and if for each α, the restriction $p_\alpha\big|_{V_\alpha}$ is the identity map on V_α. Because of the uniqueness, this definition of internal direct product is consistent with the earlier one when there are only finitely V_α's.

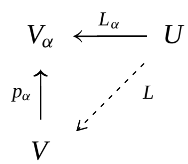

FIGURE 2.4. Universal mapping property of a direct product of vector spaces.

Proposition 2.32. Let A be a nonempty set of vector spaces over $\mathbb{F}$, and let V_α be the vector space corresponding to the member α of A. If $(V, \{p_\alpha\})$ and $(V^*, \{p^*_\alpha\})$ are two direct products of the V_α's, then the linear maps $p_\alpha : V \to V_\alpha$ and $p^*_\alpha : V^* \to V_\alpha$ are onto V_α, there exists a unique linear map $L : V^* \to V$ such that $p^*_\alpha = p_\alpha L$ for all $\alpha \in A$, and L is invertible.

PROOF. In Figure 2.4 let $U = V^*$ and $L_\alpha = p^*_\alpha$. If $L : V^* \to V$ is the linear map produced by the fact that V is a direct product, then we have $p_\alpha L = p^*_\alpha$ for all α. Reversing the roles of V and V^*, we obtain a linear map $L^* : V \to V^*$ with $p^*_\alpha L^* = p_\alpha$ for all α. Therefore $p_\alpha(LL^*) = (p_\alpha L)L^* = p^*_\alpha L^* = p_\alpha$.

In Figure 2.4 we next let $U = V$ and $L_\alpha = p_\alpha$ for all α. Then the identity 1_V on V has the same property $p_\alpha 1_V = p_\alpha$ relative to all p_α that LL^* has, and the uniqueness says that $LL^* = 1_V$. Reversing the roles of V and V^*, we obtain $L^*L = 1_{V^*}$. Therefore L is invertible.

For uniqueness suppose that $\Phi : V^* \to V$ is another linear map with $p^*_\alpha = p_\alpha \Phi$ for all $\alpha \in A$. Then the argument of the previous paragraph shows that $L^*\Phi = 1_{V^*}$. Applying L on the left gives $\Phi = (LL^*)\Phi = L(L^*\Phi) = L1_{V^*} = L$. Thus $\Phi = L$.

Finally we have to show that the α^{th} map of a direct product is onto V_α. It is enough to show that p^*_α is onto V_α. Taking V as the external direct product $\prod_{\alpha \in A} V_\alpha$ with p_α equal to the coordinate mapping, form the invertible linear map $L^* : V \to V^*$ that has just been proved to exist. This satisfies $p_\alpha = p^*_\alpha L^*$ for all $\alpha \in A$. Since p_α is onto V_α, p^*_α must be onto V_α. □

A **direct sum** of a set of vector spaces V_α over $\mathbb{F}$ for $\alpha \in A$ consists of a vector space V and a system of linear maps $i_\alpha : V_\alpha \to V$ with the following **universal mapping property**: whenever W is a vector space and $\{M_\alpha\}$ is a system of linear maps $M_\alpha : V_\alpha \to W$, then there exists a unique linear map $M : V \to W$ such that $Mi_\alpha = M_\alpha$ for all α. See Figure 2.5. The external direct sum establishes existence of a direct sum, and Proposition 2.33 below establishes its uniqueness up to isomorphism of the V's that respects the i_α's. A direct sum is said to be **internal** if each V_α is a vector subspace of V and if for each α, the map i_α is the inclusion map of V_α into V. Because of the uniqueness, this definition of internal direct sum is consistent with the earlier one when there are only finitely V_α's.

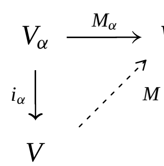

FIGURE 2.5. Universal mapping property of a direct sum of vector spaces.

Proposition 2.33. Let A be a nonempty set of vector spaces over $\mathbb{F}$, and let V_α be the vector space corresponding to the member α of A. If $(V, \{i_\alpha\})$ and $(V^*, \{i_\alpha^*\})$ are two direct sums of the V_α's, then the linear maps $i_\alpha : V_\alpha \to V$ and $i_\alpha^* : V_\alpha \to V^*$ are one-one, there exists a unique linear map $M : V \to V^*$ such that $i_\alpha^* = Mi_\alpha$ for all $\alpha \in A$, and M is invertible.

PROOF. In Figure 2.5 let $W = V^*$ and $M_\alpha = i_\alpha^*$. If $M : V \to V^*$ is the linear map produced by the fact that V is a direct sum, then we have $Mi_\alpha = i_\alpha^*$ for all α. Reversing the roles of V and V^*, we obtain a linear map $M^* : V^* \to V$ with $M^* i_\alpha^* = i_\alpha$ for all α. Therefore $(M^*M)i_\alpha = M^*i_\alpha^* = i_\alpha$.

In Figure 2.5 we next let $W = V$ and $M_\alpha = i_\alpha$ for all α. Then the identity 1_V on V has the same property $1_V i_\alpha = i_\alpha$ relative to all i_α that M^*M has, and the uniqueness says that $M^*M = 1_V$. Reversing the roles of V and V^*, we obtain $MM^* = 1_{V^*}$. Therefore M is invertible.

For uniqueness suppose that $\Phi : V \to V^*$ is another linear map with $i_\alpha^* = \Phi i_\alpha$ for all $\alpha \in A$. Then the argument of the previous paragraph shows that $M^*\Phi = 1_V$. Applying M on the left gives $\Phi = (MM^*)\Phi = M(M^*\Phi) = M1_V = M$. Thus $\Phi = M$.

Finally we have to show that the α^{th} map of a direct sum is one-one on V_α. It is enough to show that i_α^* is one-one on V_α. Taking V as the external direct sum $\bigoplus_{s\in S} V_\alpha$ with i_α equal to the embedding mapping, form the invertible linear map $M^* : V^* \to V$ that has just been proved to exist. This satisfies $i_\alpha = M^* i_\alpha^*$ for all $\alpha \in A$. Since i_α is one-one, i_α^* must be one-one. □

7. Determinants

A "determinant" is a certain scalar attached initially to any square matrix and ultimately to any linear map from a finite-dimensional vector space into itself.