We also now know that only square matrices are invertible.

(3.17) Proposition: An invertible matrix is necessarily square. More precisely, if $A \in \mathbb{F}^{m\times n}$, then (i) A 1-1 implies that $m \ge n$; and (ii) A onto implies that $m \le n$.

If $A \in \mathbb{F}^{n\times n}$ is invertible, then the first n columns of $[A,\ \mathrm{id}_n]$ are necessarily bound and the remaining n columns are necessarily free. Therefore, if $R := \mathrm{rrref}([A,\ \mathrm{id}_n])$, then $R = [\,\mathrm{id}_n, ?]$ and, with (3.11), necessarily $[A,\ \mathrm{id}_n] = AR = [A\,\mathrm{id}_n, A?]$, hence $? = A^{-1}$, i.e., $R = [\,\mathrm{id}_n, A^{-1}]$.

practical note: Although `MATLAB` provides the function `inv(A)` to generate the inverse of `A`, there is usually no reason to compute the inverse of a matrix, nor would you solve the linear system $A? = y$ in practice by computing $\mathrm{rref}([A, y])$ or by computing `inv(A)*y`. Rather, in `MATLAB` you would compute the solution of `A`$?$ `=y` as `A\y`. For this, `MATLAB` also uses elimination, but in a more sophisticated form, to keep rounding error effects as small as possible. In effect, the choice of pivot rows is more elaborate than we discussed above.

□

(3.18) Example: Triangular matrices There is essentially only one class of square matrices whose invertibility can be settled by inspection, namely the class of triangular matrices.

Assume that the square matrix A is **upper triangular**, meaning that $i > j \implies A(i,j) = 0$. If all its diagonal elements are nonzero, then each of its unknowns has a pivot row, hence is bound and, consequently, A is 1-1, hence, by (3.16)Theorem, it is invertible. Conversely, if some of its diagonal elements are zero, then there must be a first zero diagonal entry, say $A(i,i) = 0 \ne A(k,k)$ for $k < i$. Then, for $k < i$, row k is a pivot row for x_k, hence, when it comes time to decide whether x_i is free or bound, all rows not yet used as pivot rows do not involve x_i explicitly, and so x_i is free. Consequently, $\mathrm{null}\, A$ is nontrivial and A fails to be 1-1.

Exactly the same argument can be made in case A is **lower triangular**, meaning that $i < j \implies A(i,j) = 0$, provided you are now willing to carry out the elimination process from right to left, i.e., in the order x_n, x_{n-1}, etc., and, correspondingly, recognize a row as pivot row for x_k in case x_k is the last unknown that appears explicitly (i.e., with a nonzero coefficient) in that row.

(3.19) Proposition: A square triangular matrix is invertible if and only if all its diagonal entries are nonzero.

□

(3.20) Example: Interpolation If $V \in L(\mathbb{F}^n, X)$ and $Q \in L(X, \mathbb{F}^n)$, then QV is a linear map from $\mathbb{F}^n$ to $\mathbb{F}^n$, i.e., a square matrix, of order n. If QV is 1-1 or onto, then (3.16)Theorem tells us that QV is invertible. In particular, V is 1-1 and Q is onto, and so, for every $y \in \mathbb{F}^n$, there exists exactly one $p \in \operatorname{ran} V$ for which $Qp = y$. This is the essence of *interpolation*.

For example, take $X = \mathbb{R}^{\mathbb{R}}$, $V = [()^0, ()^1, \ldots, ()^{k-1}]$, hence $\operatorname{ran} V$ equals $\Pi_{<k}$, the collection of all polynomials of degree $< k$. Further, take $Q : X \to \mathbb{R}^k : f \mapsto (f(\tau_1), \ldots, f(\tau_k))$ for some fixed sequence $\tau_1 < \cdots < \tau_k$ of points. Then the equation

$$QV? = Qf$$

asks for the (power) coefficients of a polynomial of degree $< k$ that agrees with the function f at the k distinct points τ_j.

We investigate whether QV is 1-1 or onto, hence invertible. For this, consider the matrix QW, with the columns of $W := [w_1, \ldots, w_k]$ the so-called **Newton polynomial**s

$$w_j : t \mapsto \prod_{h<j} (t - \tau_h).$$

Observe that $(QW)(i,j) = (Qw_j)(\tau_i) = \prod_{h<j}(\tau_i - \tau_h) = 0$ if and only if $i < j$. Therefore, QW is square and lower triangular with nonzero diagonal entries, hence invertible by (3.19)Proposition, while w_j is a polynomial of exact degree $j - 1 < k$, hence $w_j = Vc_j$ for some k-vector c_j. It follows that the invertible matrix QW equals

$$QW = [Qw_1, \ldots, Qw_k] = [QVc_1, \ldots, QVc_k] = (QV)[c_1, \ldots, c_k].$$

In particular, QV is onto, hence invertible, hence also V is 1-1, therefore invertible as a linear map from $\mathbb{R}^k$ to its range, $\Pi_{<k}$. We have proved:

(3.21) Proposition: For every $f : \mathbb{R} \to \mathbb{R}$ and every k distinct points $\tau_1, \ldots, \tau_k$ in $\mathbb{R}$, there is exactly one choice of coefficient vector a for which the polynomial $[()^0, \ldots, ()^{k-1}]a$ of degree $< k$ agrees with f at these τ_j.

In particular, (i) the column map $[()^0, \ldots, ()^{k-1}] : \mathbb{R}^k \to \Pi_{<k}$ is invertible, and (ii) any polynomial of degree $< k$ with more than $k - 1$ distinct zeros must be 0.

□

3.17 For each of the following matrices A, use elimination (to the extent necessary) to (a) determine whether it is invertible and, if it is, to (b) construct the inverse (see the remark following (3.17)Proposition).

(a) $\begin{bmatrix} 1 & 2 & 3 \\ 2 & 3 & 4 \end{bmatrix}$; (b) $\begin{bmatrix} 1 & 2 \\ 2 & 3 \\ 3 & 4 \end{bmatrix}$; (c) $\begin{bmatrix} 1 & 2 & 3 \\ 2 & 3 & 4 \\ 3 & 4 & 5 \end{bmatrix}$; (d) $\begin{bmatrix} 1 & 2 & 3 \\ 2 & 3 & 4 \\ 3 & 4 & 4 \end{bmatrix}$; (e) $\begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 4 \\ 1 & 3 & 8 \end{bmatrix}$; (f) $[e_1 - e_3, e_2, e_3 + e_4, e_4] \in \mathbb{R}^{4\times 4}$.

3.18 (a) Construct the unique element of $\operatorname{ran}[()^0, ()^2, ()^4]$ that agrees with $()^1$ at the three points 0, 1, 2.

(b) Could (a) have been carried out if the pointset had been -1, 0, 1 (instead of 0, 1, 2)?

3.19 T/F

(a) $\begin{bmatrix} 1 & 0 & 1 \\ 0 & 2 & 0 \\ 0 & 0 & 0 \end{bmatrix}$ is in row echelon form.

(b) If all unknowns in the linear system $A? = 0$ are free, then $A = 0$;

(c) If all unknowns in the linear system $A? = 0$ are bound, then A is invertible.

(d) If some unknowns in the linear system $A? = 0$ are free, then A cannot be invertible.

(e) The inverse of an upper triangular matrix is lower triangular.

(f) A linear system of n equations in $n + 1$ unknowns always has solutions.

(g) Any square matrix in row echelon form is upper triangular.

(h) If A and B are square matrices of the same order, then $AB? = 0$ has the same number of bound unknowns as does $BA? = 0$.

(i) If A and B are square matrices of the same order, and AB is invertible, then also BA is invertible.

(j) If $\operatorname{null} A = \operatorname{null} B$, then $A? = 0$ and $B? = 0$ have the same free and bound unknowns.

4. The dimension of a vector space

Bases

The only vector spaces in which we can carry out calculations are the coordinate spaces $\mathbb{F}^n$. To calculate with other vector spaces, we have to relate them first to some coordinate space. This is true even when X is a proper subspace of $\mathbb{F}^n$, e.g., the nullspace of some matrix.

For example, we do not really compute with polynomials, we usually compute with the coefficients of the polynomial. Precisely (see (3.21)Proposition), one sets up the invertible linear map

$$\mathbb{F}^n \to \Pi_{<n} : a \mapsto a_1 + a_2 t + a_3 t^2 + \cdots + a_n t^{n-1}$$

where I have, temporarily, followed the (ancient and sometimes confusing) custom of describing the *monomials* by the list of symbols $(\ , t, t^2, t^3, \ldots)$ rather than by the nonstandard symbols $()^j$, $j = 0, 1, 2, 3, \ldots$ introduced earlier. One adds polynomials by adding their coefficients, or evaluates polynomials from their coefficients, etc. You may be so used to that, that you haven't even noticed until now that you do not work with the polynomials themselves, but only with their coefficients.

It is therefore a practically important goal to provide ways of **representing** the elements of a given vector space X by n-vectors. We do this by using linear maps from some $\mathbb{F}^n$ that have X as their range, i.e., we look for sequences $v_1, v_2, \ldots, v_n$ in X for which the linear map $[v_1, v_2, \ldots, v_n] : \mathbb{F}^n \to X$ is onto. If there is such a map for some n, then we call X **finitely generated**.

Among such onto maps $V \in L(\mathbb{F}^n, X)$, those that are also 1-1, hence invertible, are surely the most desirable ones since, for such V, there is, for any $x \in X$, exactly one $a \in \mathbb{F}^n$ with $x = Va$. Any *invertible* column map to X is, by definition, a **basis** for X.

Since $\mathrm{id}_n \in L(\mathbb{F}^n)$ is trivially invertible, it is a basis for $\mathbb{F}^n$. It is called the **natural basis for $\mathbb{F}^n$**.

The bound part, $A(:, \texttt{bound})$, of $A \in \mathbb{F}^{m\times n}$ is a basis for $\operatorname{ran} A$. You also know (from pages 36ff) how to construct a basis for the nullspace of any $A \in \mathbb{F}^{m\times n}$ from its $\operatorname{rrref}(A)$.

Here is a small difficulty with this (and any other) definition of dimension: What is the dimension of the **trivial space**, i.e., the vector space that consists of the zero vector alone? It is a perfectly well-behaved vector space (though a bit limited, – except as a challenge to textbook authors when it comes to discussing its basis).

We deal with it here by considering $V \in L(\mathbb{F}^n, X)$ even when $n = 0$. Since $\mathbb{F}^n$ consists of lists of n items (each item an element from $\mathbb{F}$), the peculiar space $\mathbb{F}^0$ must consist of lists of *no* items, i.e., of *empty* lists. There is only one empty list (of scalars), hence $\mathbb{F}^0$ has just one element, the empty list, $(\,)$, and this element is necessarily the neutral element (or, zero vector) for this space. Correspondingly, there is exactly one *linear* map from $\mathbb{F}^0$ into X, namely the map $\mathbb{F}^0 \to X : () = 0 \mapsto 0$. Since this is a linear map from $\mathbb{F}^0$, we call it the column map into X with *no* columns, and denote it by $[\,]$. Thus,

$$[\,] : \mathbb{F}^0 \to X : () = 0 \mapsto 0. \tag{4.1}$$

Note that $[\,]$ *is 1-1.* Note also that the range of $[\,]$ consists of the trivial subspace, $\{0\}$. In particular, the column map $[\,]$ is *onto* $\{0\}$, hence is invertible, as map from $\mathbb{F}^0$ to $\{0\}$. It follows that $[\,]$ is a basis for $\{0\}$. Isn't Mathematics wonderful! - As it turns out, the column map $[\,]$ will also be very helpful below.

Here are some standard terms related to bases of a vector space:

Definition: The range of $V := [v_1, v_2, \ldots, v_n]$ is called the **span of the sequence** $v_1, v_2, \ldots, v_n$:

$$\operatorname{span}(v_1, v_2, \ldots, v_n) := \operatorname{ran} V.$$

$x \in X$ is said to be **linearly dependent on** $v_1, v_2, \ldots, v_n$ in case $x \in \operatorname{ran} V$, i.e., in case x is a **linear combination of the** v_j. Otherwise x is said to be **linearly independent of** $v_1, v_2, \ldots, v_n$.

$v_1, v_2, \ldots, v_n$ is said to be **linearly independent** in case V is 1-1, i.e., in case $Va = 0$ implies $a = 0$ (i.e., the only way to write the zero vector as a linear combination of the v_j is to choose all the weights equal to 0).

$v_1, v_2, \ldots, v_n$ is said to be **spanning for** X in case V is onto, i.e., in case $\operatorname{span}(v_1, v_2, \ldots, v_n) = X$.

$v_1, v_2, \ldots, v_n$ is said to be a **basis for** X in case V is invertible, i.e., 1-1 and onto.

If V is invertible, then $V^{-1}x$ is an n-vector, called the **coordinate vector for** x **with respect to the basis** $v_1, v_2, \ldots, v_n$.

You may wonder why there are all these terms in use for the *sequence* $v_1, v_2, \ldots, v_n$, particularly when the corresponding terms for the *map* V are so much shorter and to the point. I don't know the answer. However, bear in mind that the terms commonly used are those for sequences.

A major use of the basis concept is the following which generalizes the way we earlier constructed arbitrary linear maps from $\mathbb{F}^n$.

(4.2) Proposition: Let $V = [v_1, v_2, \ldots, v_n]$ be a basis for the vector space X, and let Y be an arbitrary vector space. Any map $f : \{v_1, \ldots, v_n\} \to Y$ has exactly one extension to a linear map A from X to Y. In other words, we can choose the values of a linear map on the columns of a basis arbitrarily and, once chosen, this pins down the linear map everywhere.

Proof: The map $A := [f(v_1), \ldots, f(v_n)]V^{-1}$ is linear, from X to Y, and carries v_j to $f(v_j)$ since $V^{-1}v_j = e_j$, all j. This shows existence. Further, if also $B \in L(X, Y)$ with $Bv_j = f(v_j)$, all j, then $BV = [f(v_1), \ldots, f(v_n)] = AV$, therefore $B = A$ (since V is invertible). $\square$

4.1 Describe what the $n \times n$-matrix $A = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 1 & \cdots & 0 & 0 \\ \cdot & \cdot & \cdot & \cdots & \cdot & \cdot \\ 0 & 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & 0 & \cdots & 0 & 1 \\ 0 & 0 & 0 & \cdots & 0 & 0 \end{bmatrix}$ does to all the vectors e_j, i.e., give a simple formula for Ae_j. Deduce from your formula that $\operatorname{ran} A^n = \{0\}$, hence that $A^n = 0$.

4.2 Prove: $A \in L(X)$ commutes with every $B \in L(X)$ if and only if $A = \alpha \operatorname{id}_X$, i.e., a scalar multiple of the identity.

4.3 Let $X \times Y$ be the product space of the vector spaces X and Y. The map $f : X \times Y \to \mathbb{F}$ is **bilinear** if it is linear in each slot, i.e., if $f(\cdot, y) \in L(X, \mathbb{F})$ for all $y \in Y$, and $f(x, \cdot) \in L(Y, \mathbb{F})$ for every $x \in X$.

(i) Prove that, for every $A \in \mathbb{F}^{m \times n}$, the map $f_A : \mathbb{F}^m \times \mathbb{F}^n : (x, y) \mapsto y^t A x$ is bilinear.

(ii) Prove that, for every bilinear $f : \mathbb{F}^m \times \mathbb{F}^n \to \mathbb{F}$, there exists exactly one $A \in \mathbb{F}^{m \times n}$ with $f_A = f$.

(iii) Prove that the map $A \mapsto f_A$ is an invertible linear map on $\mathbb{F}^{m \times n}$ to the vector space $BL(\mathbb{F}^m, \mathbb{F}^n)$ of all bilinear maps on $\mathbb{F}^m \times \mathbb{F}^n$ under pointwise vector operations.

4.4 `MATLAB`'s command `yy = interp1(x,y,xx,'spline')` returns the value(s) at `xx` of a certain function f that matches the data given by `x`, `y`, in the sense that f(`x(i)`) = `y(i)` for `i=1:n`, with n the length of both `x` and `y` (and assuming that the entries of `x` are pairwise distinct). (If you wanted to look at f on the interval $[a\,..\,b]$, you might choose `xx = linspace(a,b,N+1);` with `N` some suitably large number, and then `plot(xx,yy)`.)

(a) Generate some numerical evidence for the claim that (up to roundoff) the map `y`$\mapsto f$ provided by this command is linear.

(b) Assuming that the map is linear, deduce from the above description of the map that it must be 1-1, hence a basis for its range.

(c) Still assuming that the map $y \mapsto f$ provided by that command is indeed linear, hence a column map, provide a plot of each of its columns, as functions on the interval $[0\,..\,3]$, for the specific choice `0:3` for `x`.

(d) (quite open-ended) Determine as much as you can about the elements of the range of this column map.

(e) Is the map still linear if you replace `'spline'` by `'cubic'`?

Construction of a basis

Next, we consider the construction of a basis. This can be done either by *extending a 1-1 column map* V to a basis, or by *thinning an onto column map* W to a basis. For this, remember that, for two column maps V and W into some vector space X, we agreed to mean by $V \subset W$ that V can be obtained from W by thinning, i.e., by omitting zero or more columns from W, and W can be obtained from V by extending, i.e., by inserting zero or more columns.

In the discussion to follow, it is convenient to classify the columns of a column map as *bound* or *free*, using (3.5)Corollary as a guide. Specifically, we call a column **free** if it is a weighted sum of the columns to its left; otherwise, we call it **bound**.

For example, if $V \subset W$, then any free column of V is also free as a column of W, while a bound column of V may possibly be free as a column of W unless $W = [V, U]$.

(4.3) Lemma: The kth column of the column map V is free if and only if $\operatorname{null} V$ contains a vector whose last nonzero entry is its kth.

Proof: The kth column of $V = [v_1, \ldots, v_n] \in L(\mathbb{F}^n, X)$ is free iff $v_k \in \operatorname{ran}[v_1, \ldots, v_{k-1}]$. In particular, the first column is free iff it is 0 (recall that $\operatorname{ran}[\,] = \{0\}$).

If the kth column is free, then $v_k = [v_1, \ldots, v_{k-1}]a$ for some $a \in \mathbb{F}^{k-1}$, hence $(a, -1, 0, \ldots, 0) \in \mathbb{F}^n$ is a vector in $\operatorname{null} V$ whose last nonzero entry is its kth. Conversely if $x \in \operatorname{null} V$ with $x_k \neq 0 = x_{k+1} = \cdots = x_n$, then $[v_1, \ldots, v_{k-1}]x_{1:k-1} + v_k x_k = 0$, therefore, as $x_k \neq 0$, $v_k = [v_1, \ldots, v_{k-1}](x_{1:k-1}/(-x_k))$ showing that the kth column is free. □

(4.4) Corollary: A column map is 1-1 if and only if all of its columns are bound.

We are ready for the following algorithm which extracts from any column map W a basis for its range.

(4.5) Basis Selection Algorithm:
input: the column map W
$V \leftarrow [\,]$;
for $w \in W$ **do:**
    if $w \notin \operatorname{ran} V$**, then** $V \leftarrow [V, w]$**; endif**
enddo
output: the column map V

Proposition: The output of the Basis Selection Algorithm is a basis for the range of its input.

Proof: The resulting V has the same range as W (since the only columns of W not explicitly columns of V are those that are already in the range of V). In addition, by construction, every column of V is bound, hence V is 1-1 by (4.4)Corollary, therefore a basis for its range. □

(4.6) Proposition: Any onto column map can be thinned to a basis.

Now note that the Basis Selection Algorithm will put any bound column of W into the resulting basis, V. In particular, if $W = [U, Z]$ with U 1-1, then, as already remarked just prior to (4.3)Lemma, all columns of U will be bound also as columns of W, hence will end up in the resulting basis. This proves

(4.7) Proposition: Any 1-1 column map into a finitely generated vector space can be extended to a basis for that space.

If V is a 1-1 column map into X then, by (4.4)Corollary, all its columns are bound. Hence if V is **maximally 1-1** into X, meaning that $[V, w]$ fails to be 1-1 for every $w \in X$, then that additional column must be free, i.e., $w \in \operatorname{ran} V$ for all $w \in X$, showing that then V is also onto, hence a basis. This proves

(4.8) Corollary: Any maximally 1-1 column map into a vector space is a basis for that space.

If W is a column map onto X, then, by (4.6), it can always be thinned to a basis. Hence, if W is **minimally onto**, meaning that no $V \subset W$ (other than W) is onto, then W itself must be that basis.

(4.9) Corollary: Any minimally onto column map into a vector space is a basis for that space.

4.5 How would you carry out the (4.5) Basis Selection Algorithm for the special case that W is a matrix? (Hint: (3.2)).

4.6 Try out your answer to the previous problem on the specific matrix $W = \begin{bmatrix} 0 & 2 & 0 & 2 & 5 & 4 & 0 & 6 \\ 0 & 1 & 0 & 1 & 2 & 2 & 0 & 3 \\ 0 & 2 & 0 & 2 & 5 & 4 & -1 & 7 \end{bmatrix}$.

Dimension

(4.10) Lemma: Any two bases for a vector space have the same number of columns.

This number of columns in any basis for X is denoted

$$\dim X$$

and is called the **dimension of** X.

Proof: Let $V = [v_1, \ldots, v_n]$ and W be bases for the vector space X. Since W is onto, we have $v_n = Wa$ for some a, hence $(-1, a)$ is a nonzero vector in $\operatorname{null}[v_n, W]$ and so, $[v_n, W]$ is onto but not 1-1, hence, by (4.4)Corollary, has some free columns. Yet, since $v_n \neq 0$, all these free columns must be in W and, dropping any one of them, say the first one, we get the onto column map $[v_n, W_1]$.

It follows that $v_{n-1} = [v_n, W_1]a$ for some vector a, hence the nonzero vector $(-1, a)$ is in $\operatorname{null}[v_{n-1}, v_n, W_1]$, and so, $[v_{n-1}, v_n, W_1]$ is onto but not 1-1, therefore, by (4.4)Corollary, it must have some free columns. Yet, since the first two columns are part of a 1-1 map, all these free columns must be in W_1 and, dropping any one of them, say the first one, we get the onto column map $[v_{n-1}, v_n, W_2]$.

You get the pattern: as we introduce in this way the columns of V one by one, there must always be at least one column from W still to drop. In other words, we must have $\#V \le \#W$. However, since also W is 1-1 and V is onto, we also must have $\#W \le \#V$, and that finishes the proof.

Of course, if you are willing to make use of a result from the previous chapter, then the proof of this lemma is immediate: Let $V \in L(\mathbb{F}^n, X)$ and $W \in L(\mathbb{F}^m, X)$ be bases for X. Then, $W^{-1}V$ is an invertible linear map from $\mathbb{F}^n$ to $\mathbb{F}^m$, hence an invertible matrix and therefore, by (3.17)Proposition(i), necessarily a square matrix, i.e., $n = m$. □

Notice that we have actually proved the stronger statement

(4.11) Lemma: If V and W are column maps into X, and V is 1-1 and W is onto, then $\#V \le \#W$.

Again, also this stronger result is an immediate consequence of something proved in the previous chapter: Since W is onto, each column v_j of V can be written as $v_j = Wc_j$ for some vector c_j. Hence $V = WC$ for some matrix C and, since V is 1-1, so must C be. By (3.17)Proposition(i) or its antecedent, (3.6)Theorem, this implies that C cannot have more columns than rows, i.e., $\#V = \#C \le \dim \operatorname{tar} C = \dim \operatorname{dom} W = \#W$.

Since id_n is a basis for $\mathbb{F}^n$ and has n columns, we conclude that the *n-dimensional coordinate space has, indeed, dimension n*. In effect, $\mathbb{F}^n$ is the prototypical vector space of dimension n. Any n-dimensional vector space X is connected to $\mathbb{F}^n$ by invertible linear maps, the bases for X.

Note that the trivial vector space, $\{0\}$, has dimension 0 since its (unique) basis has no columns.

(4.12) Example: The dimension of $\Pi_k(\mathbb{R}^d)$. The space $\Pi_k(\mathbb{R}^d)$ of **d-variate polynomials of degree** $\le k$ is, by definition, the range of the column map $V := [()^\alpha : |\alpha| \le k]$, with

$$()^\alpha : \mathbb{R}^d \to \mathbb{R} : t \mapsto t^\alpha := t_1^{\alpha_1} \cdots t_d^{\alpha_d}$$

a nonstandard notation for the α-power function, with $\alpha \in \mathbb{Z}_+^d$, i.e., α any d-vector with nonnegative integer entries, and with $|\alpha| := \sum_j \alpha_j$. For $d = 1$, it is the space of univariate polynomials of degree $\le k$, and we showed its dimension to be $k+1$ by showing that, in that case, V is 1-1.

When $d = 1$, then V can be seen to be 1-1 also by considering the 'data map'

$$Q : \Pi_k \to \mathbb{R}^{k+1} : p \mapsto (p(0), Dp(0), D^2p(0)/2, \ldots, D^k p(0)/k!),$$

for which we have $QV = \mathrm{id}$, hence V is 1-1.

An analogous argument, involving the 'data map'

$$p \mapsto (D^\alpha p(0)/\alpha! : \alpha \in \mathbb{Z}_+^d, |\alpha| \le k),$$

with $\alpha! := \alpha_1! \cdots \alpha_d!$, shows that

$$\dim \Pi_k(\mathbb{R}^d) = \#\{\alpha \in \mathbb{Z}_+^d : |\alpha| \le k\},$$

and the latter number can be shown to equal $\binom{k+d}{d}$. □

4.7 Prove that the space $\Pi_2(\mathbb{R}^2)$ of bivariate polynomials of total degree ≤ 2 has dimension 6.

4.8 Prove that a vector space of dimension n has subspaces of dimension j for each $j = 0{:}n$.

Some uses of the dimension concept

Here is a major use of the dimension concept as it relates to *vector spaces*.

(4.13) Proposition: If X, Y are vector spaces with $X \subset Y$ and $\dim Y < \infty$, then $\dim X \le \dim Y$, with equality iff $X = Y$.

Proof: Since there is *some* 1-1 column map into X (e.g., the unique linear map from $\mathbb{F}^0$ into X), while $\dim Y$ is an upper bound on the number of columns in any 1-1 column map into $X \subset Y$ (by (4.7)Proposition), there exists a maximally 1-1 column map V into X. By (4.8)Corollary, any such V is necessarily a basis for X, hence X is finitely generated. By (4.7)Proposition, we can extend V to a basis $[V, W]$ for Y. Hence, $\dim X \le \dim Y$ with equality iff $W = [\,]$, i.e., iff $X = Y$. □

Note the following important (nontrivial) part of (4.13)Proposition:

(4.14) Corollary: Any linear subspace of a finite-dimensional vector space is finite-dimensional.

The dimension concept is usually applied to *linear maps* by way of the following formula.

(4.15) Dimension Formula: For any linear map A with finite-dimensional domain,

$$\dim \operatorname{dom} A = \dim \operatorname{ran} A + \dim \operatorname{null} A.$$

Proof: Since $\operatorname{dom} A$ is finite-dimensional, so is $\operatorname{null} A$ (by (4.14)Corollary), hence $\operatorname{null} A$ has a basis, $V \in L(\mathbb{F}^n, \operatorname{null} A)$ say. By (4.7)Proposition, we can extend this to a basis $[V, U]$ for $\operatorname{dom} A$. Let $r := \#U$. Then, $[V, U]$ is invertible and $\dim \operatorname{dom} A - \dim \operatorname{null} A = (n + r) - n = r$.

It remains to prove that $\dim \operatorname{ran} A = r$. For this, we prove that $AU : \mathbb{F}^r \to \operatorname{ran} A$ is invertible.

Since $A[V, U] = [AV, AU]$ maps onto $\operatorname{ran} A$ and $AV = 0$, already AU must map onto $\operatorname{ran} A$, i.e., AU is onto.

Moreover, AU is 1-1: For, if $AUa = 0$, then $Ua \in \operatorname{null} A$, hence, since V maps onto $\operatorname{null} A$, there is some b so that $Ua = Vb$. This implies that $[V, U](b, -a) = 0$ and, since $[V, U]$ is 1-1, this shows that, in particular, $a = 0$. □

is defined,

(4.16) Corollary: Let $A \in L(X, Y)$.
(i) If $\dim X < \dim Y$, then A cannot be onto.
(ii) If $\dim X > \dim Y$, then A cannot be 1-1.
(iii) If $\dim X = \dim Y < \infty$, then A is onto if and only if A is 1-1. (This implies (2.18)!)

Proof: (i) $\dim \operatorname{ran} A \le \dim \operatorname{dom} A = \dim X < \dim Y = \dim \operatorname{tar} A$, hence $\operatorname{ran} A \ne \operatorname{tar} A$.

(ii) $\dim \operatorname{null} A = \dim \operatorname{dom} A - \dim \operatorname{ran} A = \dim X - \dim \operatorname{ran} A \ge \dim X - \dim Y > 0$, hence $\operatorname{null} A \ne \{0\}$.

(iii) If $\dim X = \dim Y$, then $\dim \operatorname{tar} A = \dim \operatorname{dom} A = \dim \operatorname{ran} A + \dim \operatorname{null} A$, hence A is onto (i.e., $\operatorname{tar} A = \operatorname{ran} A$) if and only if $\dim \operatorname{null} A = 0$, i.e., A is 1-1. □

(4.17) Lemma: Let X, Y be vector spaces, and assume that X is finite-dimensional. Then $\dim X = \dim Y$ if and only if there exists an invertible $A \in L(X, Y)$.

Proof: Let $n := \dim X$. Since $n < \infty$, there exists an invertible $V \in L(\mathbb{F}^n, X)$ (, a basis for X). If now $A \in L(X, Y)$ is invertible, then AV is an invertible linear map from $\mathbb{F}^n$ to Y, hence $\dim Y = n = \dim X$. Conversely, if $\dim Y = \dim X$, then there exists an invertible $W \in L(\mathbb{F}^n, Y)$; but then WV^{-1} is an invertible linear map from X to Y. □

For the next general result concerning the dimension concept, recall that both the sum

$$Y + Z := \{y + z : y \in Y, z \in Z\}$$

and the intersection $Y \cap Z$ of two linear subspaces is again a linear subspace.

(4.18) Proposition: If Y and Z are linear subspaces of the finite-dimensional vector space X, then

$$\dim(Y + Z) = \dim Y + \dim Z - \dim(Y \cap Z). \tag{4.19}$$

Proof 1: $Y \cap Z$ is a linear subspace of X, hence is finite-dimensional (by (4.14)Corollary), hence $Y \cap Z$ has a basis, V say. Extend it, as we may (by (4.7)Proposition), to a basis $[U, V]$ of Y and to a basis $[V, W]$ of Z, and consider the column map $[U, V, W]$.

We claim that $[U, V, W]$ is 1-1. Indeed, if $[U, V, W](a, b, c) = 0$, then $[U, V](a, b) = -Wc$, with the left side in Y and the right side in Z, hence both are in $Y \cap Z = \operatorname{ran} V$. Therefore, $-Wc = Vd$ for some d, hence $[V, W](d, c) = 0$, and as $[V, W]$ is 1-1, it follows, in particular, that $c = 0$. This leaves $[U, V](a, b) = 0$ and, since $[U, V]$ is 1-1 by construction, now also $(a, b) = 0$.

We conclude that $[U, V, W]$ is a basis for its range, and that range is $\operatorname{ran}[U, V, W] = \operatorname{ran}[U, V, V, W] = \operatorname{ran}[U, V] + \operatorname{ran}[V, W] = Y + Z$. Therefore, $\dim(Y + Z) = \#U + \#V + \#W = \#[U, V] + \#[V, W] - \#V = \dim Y + \dim Z - \dim(Y \cap Z)$. □

Proof 2: The following alternative proof shows (4.19) to be a special case of the (4.15)Dimension Formula, and provides a way to construct a basis for $Y \cap Z$ from bases for Y and Z.

Consider the column map $A := [U, W]$ with U a basis for Y and W a basis for Z. Since $\dim \operatorname{dom} A = \#U + \#W = \dim Y + \dim Z$ and $\operatorname{ran} A = Y + Z$, the formula (4.19) follows from the (4.15)Dimension Formula, once we show that $\dim \operatorname{null} A = \dim Y \cap Z$. For this, let $x \in Y \cap Z$. Then $x = Ua = Wb$ for some a and b, therefore $A(a, -b) = [U, W](a, -b) = Ua - Wb = x - x = 0$, hence $(a, -b) \in \operatorname{null} A$. Hence, $(a, -b) = Cc$ for some c and with $C =: [C_U; C_W]$ a basis for $\operatorname{null} A$. In particular, $a = C_U c$, showing that the column map UC_U has all of $Y \cap Z$ in its range. On the other hand, $0 = AC = UC_U + WC_W$, hence $UC_U = -WC_W$ and, in particular, UC_U maps into $Y \cap Z$, hence onto $Y \cap Z$. Finally, UC_U is 1-1: for, if $UC_U a = 0$, then $C_U a = 0$ since U is 1-1, but then also $WC_W a = -UC_U a = 0$, hence also $C_W a = 0$, therefore $Ca = 0$ and so $a = 0$ since C is 1-1 by assumption. Altogether, this shows that UC_U is a basis for $Y \cap Z$, hence $\dim Y \cap Z = \#UC_U = \#C = \dim \operatorname{null} A$. □

Here are three of several corollaries of this basic proposition to be used in the sequel.

(4.20) Corollary: If $[V, W]$ is 1-1, then $\operatorname{ran} V \cap \operatorname{ran} W$ is trivial.

(4.21) Corollary: If $\dim Y + \dim Z > \dim X$ for some linear subspaces Y and Z of the finite-dimensional vector space X, then $Y \cap Z$ is a *nontrivial* linear subspace, i.e., $Y \cap Z$ contains nonzero elements.

(4.22) Corollary: If Y and Z are linear subspaces of the finite-dimensional vector space X, and $Y \cap Z = \{0\}$, then

$$\dim Y + \dim Z \le \dim X,$$

with equality if and only if $X = Y + Z$.

4.9 For each of the following linear maps, determine its range and its nullspace. Make as much use of the Dimension Formula as possible. (You may assume, if need be, that $V_k := [()^0, ()^1, \ldots, ()^k]$ is a basis for Π_k since it is proved in the Notes). (a) $D : \Pi_k \to \Pi_{k-1} : p \mapsto Dp$, with Dp the first derivative of p. (b) $I : \Pi_{k-1} \to \Pi_k : p \mapsto \int_0^{\cdot} p(s)ds$, i.e., Ip is the primitive or antiderivative of p that vanishes at 0, i.e., $(Ip)(t) = \int_0^t p(s)ds$. (c) $A : \Pi_k \to \Pi_k : p \mapsto Dp + p$.

4.10 Prove that $V := [()^0, ()^1, ()^2 - 1, 4()^3 - 3()^1, 8()^4 - 8()^2 + 1]$ is a basis for Π_4.

4.11 Call $(Y_0, \ldots, Y_r)$ a **proper chain** in the vector space X if each Y_j is a subspace and $Y_0 \subsetneq Y_2 \subsetneq \cdots \subsetneq Y_r$. Prove that, for any such proper chain, $r \le \dim X$, with equality if and only if $\dim Y_j = j$, $j = 0{:}\dim X$.

4.12 Let d be any scalar-valued map, defined on the collection of all linear subspaces of a finite-dimensional vector space X, that satisfies the following two conditions: (i) $Y \cap Z = \{0\} \implies d(Y + Z) = d(Y) + d(Z)$; (ii) $\dim Y = 1 \implies d(Y) = 1$.

Prove that $d(Y) = \dim Y$ for every linear subspace of X.

4.13 Prove: for any $A \in L(X, Y)$ and any linear subspace Z of X, $\dim A(Z) = \dim Z - \dim(Z \cap (\operatorname{null} A))$.

4.14 The **defect** of a linear map is the dimension of its nullspace: $\operatorname{defect}(A) := \dim \operatorname{null} A$. (a) Prove that $\operatorname{defect}(B) \le \operatorname{defect}(AB) \le \operatorname{defect}(A) + \operatorname{defect}(B)$. (b) Prove: If $\dim \operatorname{dom} B = \dim \operatorname{dom} A$, then also $\operatorname{defect}(A) \le \operatorname{defect}(AB)$. (c) Give an example of linear maps A and B for which AB is defined and for which $\operatorname{defect}(A) > \operatorname{defect}(AB)$.

4.15 Let $A \in L(X, Y)$, $B \in L(X, Z)$, with Y finite-dimensional. There exists $C \in L(Y, Z)$ with $A = CB$ if and only if $\operatorname{null} B \subset \operatorname{null} A$.

4.16 Prove: Assuming that the product ABC of three linear maps is defined, $\dim \operatorname{ran}(AB) + \dim \operatorname{ran}(BC) \le \dim \operatorname{ran} B + \dim \operatorname{ran}(ABC)$.

4.17 Factor space: Let Y be a linear subspace of the vector space X and consider the collection

$$X/Y := \{x + Y : x \in X\}$$

of subsets of X, with

$$x + Y := \{x\} + Y = \{x + y : y \in Y\}.$$

(i) Prove that the map

$$f : X \to X/Y : x \mapsto x + Y$$

is linear with respect to the addition

$$M + N := \{m + n : m \in M, n \in N\}$$

and the multiplication by a scalar

$$\alpha M := \begin{cases} \{\alpha m : m \in M\}, & \text{if } \alpha \ne 0; \\ Y, & \text{if } \alpha = 0, \end{cases}$$

and has Y as its nullspace.

(ii) Prove that, with these vector operations, X/Y is a linear space. (X/Y is called a **factor space**.)

(iii) Prove that $\dim X/Y = \operatorname{codim} Y$.

The dimension of $\mathbb{F}^T$

Recall from (2.2) that $\mathbb{F}^T$ is the set of all scalar-valued maps on the set T, with the set T, offhand, arbitrary.

The best known instance is n-dimensional coordinate space

$$\mathbb{F}^n := \mathbb{F}^{\underline{n}},$$

with $T = \underline{n} := \{1, 2, \ldots, n\}$. The vector space $\mathbb{F}^{m\times n}$ of all $(m \times n)$-matrices is another instance; here $T = \underline{m} \times \underline{n} := \{(i,j) : i = 1{:}m; j = 1{:}n\}$.

(4.23) Proposition: If $\#T :=$ number of elements of T is finite, then $\dim \mathbb{F}^T = \#T$.

Proof: Since T is finite, $\#T =: n$ say, we can order its elements, i.e., there is an invertible map $s : \underline{n} \to T$ (in fact, there are $n! = 1 \cdot 2 \cdots n$ such). This induces the map

$$V : \mathbb{F}^n \to \mathbb{F}^T : f \mapsto f \circ s^{-1}$$

which is linear (since, in both spaces, the vector operations are pointwise), and is invertible since it has

$$\mathbb{F}^T \to \mathbb{F}^n : g \mapsto g \circ s$$

as its inverse. Hence, V is a basis for $\mathbb{F}^T$ (the **natural basis**). □

Note how we managed this without even exhibiting the columns of V. To be sure, the jth column V is the function $v_j : T \to \mathbb{F} : s_k \mapsto \delta_{kj}$ that maps s_j to 1 and maps any other $t \in T$ to 0.

Corollary: $\dim \mathbb{F}^{m\times n} = mn$.

Proof: In this case, $\mathbb{F}^{m\times n} = \mathbb{F}^T$ with $T = \underline{m} \times \underline{n} := \{(i,j) : i = 1{:}m; j = 1{:}n\}$, hence $\#T = mn$. □

(4.24) Corollary: $\dim L(X,Y) = \dim X \cdot \dim Y$.

Proof: Assuming that $n := \dim X$ and $m := \dim Y$ are finite, we can represent every $A \in L(X,Y)$ as a matrix $\widehat{A} := W^{-1}AV \in \mathbb{F}^{m\times n}$, with V a basis for X and W a basis for Y. This sets up a map

$$R : L(X,Y) \to \mathbb{F}^{m\times n} : A \mapsto \widehat{A} = W^{-1}AV,$$

and this map is linear and invertible (indeed, its inverse is the map $\mathbb{F}^{m\times n} \to L(X,Y) : B \mapsto WBV^{-1}$). Consequently, by (4.17)Lemma, $L(X,Y)$ and $\mathbb{F}^{m\times n}$ have the same dimension. □

Corollary: If $\#T \not< \infty$, then $\mathbb{F}^T$ is not finite-dimensional.

Proof: For every finite $S \subset T$, $\mathbb{F}^T$ contains the linear subspace

$$\{f \in \mathbb{F}^T : f(t) = 0, \text{all } t \notin S\}$$

of dimension equal to $\dim \mathbb{F}^S = \#S$. If $\#T \not< \infty$, then T contains finite subsets S of arbitrarily large size, hence $\mathbb{F}^T$ contains linear subspaces of arbitrarily large dimension, hence cannot itself be finite-dimensional, by (4.13)Proposition. □

4.18 Prove: *The dimension of the vector space of all upper triangular matrices of order n is $(n+1)n/2$.*

Direct sums

A very useful coarsening of the basis concept concerns the sum of subspaces.

Let $Y_1, \ldots, Y_r$ be linear subspaces of the vector space X, let V_j be a column map onto Y_j, all j, and consider the column map

$$V := [V_1, \ldots, V_r].$$

To be sure, we could have also started with some arbitrary column map V into X, arbitrarily grouped its columns to obtain $V = [V_1, \ldots, V_r]$, and then defined $Y_j := \operatorname{ran} V_j$, all j.

Either way, any $a \in \operatorname{dom} V$ is of the form $(a_1, \ldots, a_r)$ with $a_j \in \operatorname{dom} V_j$, all j. Hence

$$\operatorname{ran} V = \{V_1 a_1 + \cdots + V_r a_r : a_j \in \operatorname{dom} V_j, j = 1{:}r\} = \{y_1 + \cdots + y_r : y_j \in Y_j, j = 1{:}r\} =: Y_1 + \cdots + Y_r,$$

the *sum* of the subspaces $Y_1, \ldots, Y_r$.

Think of this sum, as you may, as the range of the map

$$A : Y_1 \times \cdots \times Y_r \to X : (y_1, \ldots, y_r) \mapsto y_1 + \cdots + y_r. \tag{4.25}$$

Having this map A onto says that every $x \in X$ can be written in the form $y_1 + \cdots + y_r$ with $y_j \in Y_j$, all j. In other words, X is the sum of the Y_j. In symbols,

$$X = Y_1 + \cdots + Y_r.$$

Having A also 1-1 says that there is *exactly one way* to write each $x \in X$ as such a sum. In this case, we write

$$X = Y_1 \dot{+} \cdots \dot{+} Y_r,$$

and say that X is the **direct sum** of the subspaces Y_j. Note the dot atop the plus sign, to indicate the special nature of this sum. Some books would use instead the encircled plus sign, $\oplus$, but we reserve that sign for an even more special direct sum in which the summands Y_j are 'orthogonal' to each other; see the chapter on inner product spaces.

(4.26) Proposition: Let V_j be a basis for the linear subspace Y_j of the vector space X, $j = 1{:}r$, and set $V := [V_1, \ldots, V_r]$. Then, the following are equivalent.

(i) $X = Y_1 \dot{+} \cdots \dot{+} Y_r$.

(ii) V is a basis for X.

(iii) $X = Y_1 + \cdots + Y_r$ and $\dim X \ge \dim Y_1 + \cdots + \dim Y_r$.

(iv) For each j, $Y_j \cap Y_{\backslash j} = \{0\}$, with $Y_{\backslash j} := Y_1 + \cdots + Y_{j-1} + Y_{j+1} + \cdots + Y_r$, and $\dim X \le \dim Y_1 + \cdots + \dim Y_r$.

Proof: Since $\operatorname{dom} V = \operatorname{dom} V_1 \times \cdots \times \operatorname{dom} V_r$, and V_j is a basis for Y_j, all j, the linear map

$$C : \operatorname{dom} V \to Y_1 \times \cdots \times Y_r : a = (a_1, \ldots, a_r) \mapsto (V_1 a_1, \ldots, V_r a_r)$$

is invertible and $V = AC$, with A as given in (4.25). Hence, V is invertible if and only if A is invertible. This proves that (i) and (ii) are equivalent.

Also, (ii) implies (iii). As to (iii) implying (ii), the first assumption of (iii) says that V is onto X, and the second assumption says that $\dim \operatorname{dom} V = \#V \le \dim X$, hence V is minimally onto and therefore a basis for X.

As to (ii) implying (iv), the first claim of (iv) is a special case of (4.20)Corollary, and the second claim is immediate.

Finally, as to (iv) implying (ii), assume that $0 = Va = \sum_j V_j a_j$. Then, for any j, $y := V_j a_j = -\sum_{i \ne j} V_i a_i \in Y_j \cap Y_{\backslash j}$, hence $y = 0$ by the first assumption and, since V_j is a basis for Y_j, hence 1-1, this implies that $a_j = 0$. In other words, V is 1-1, while, by the second assumption, $\#V = \sum_j \dim Y_j \ge \dim X$, hence V is maximally 1-1, therefore a basis for X. □

(4.27) Corollary: If V is a basis for X, then, for any grouping $V =: [V_1, \ldots, V_r]$ of the columns of V, X is the direct sum of the linear subspaces $\operatorname{ran} V_j$, $j = 1{:}r$.

One particular grouping is, of course, $V_j = [v_j]$, all j, in which case each $Y_j := \operatorname{ran} V_j$ is a one-dimensional linear subspace, i.e., a straight line through the origin, and we see $X = \operatorname{ran} V$ as the direct sum of these straight lines, each of which we are accustomed to think of as a **coordinate axis**.

This is illustrated in (4.28)Figure for the special case $\operatorname{ran} V = \mathbb{R}^2$, hence V has just two columns. We see each $x \in \mathbb{R}^2$ written as the sum $x = y_1 + y_2$, with $y_j = a_j v_j \in Y_j = \operatorname{ran}[v_j]$ the Y_j**-component** of x (and, of course, $a = (a_1, a_2)$ the coordinate vector of x with respect to the basis V).

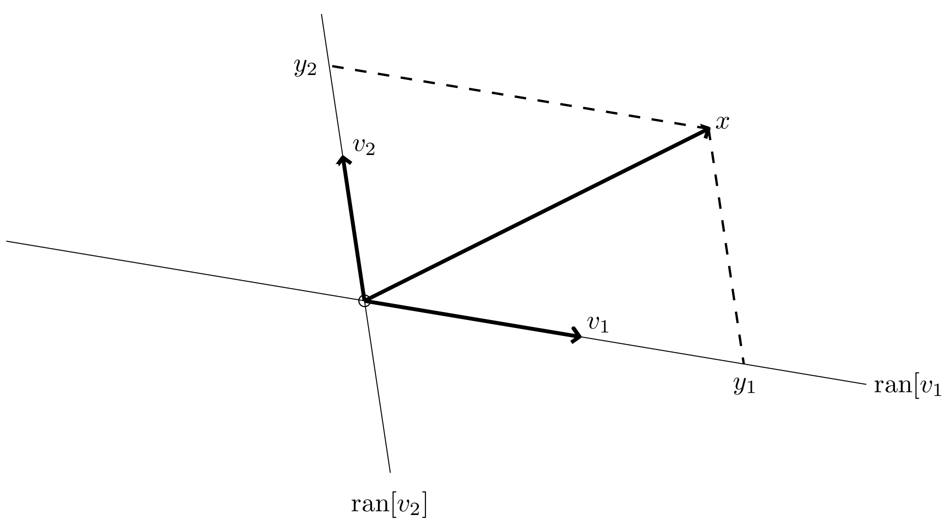

(4.28) Figure. A basis provides a coordinate system.

The direct sum construct is set up in just the same way, except that the Y_j may be planes or even higher-dimensional subspaces rather than just straight lines.

4.19 When X is the direct sum of Y and Z, then Z is said to be a **complement** of Y. With Y and Z linear subspaces of the finite-dimensional vector space X, prove the following assertions concerning complements.

(i) Y has a complement.

(ii) If both Z and Z_1 complement Y, then $\dim Z = \dim Z_1$. (This dimension is known as the **codimension** of Y, and is denoted $\operatorname{codim} Y$.) In particular, $\operatorname{codim} Y = \dim X - \dim Y$.

(iii) $\operatorname{codim}(Y + Z) = \operatorname{codim} Y + \operatorname{codim} Z - \operatorname{codim}(Y \cap Z)$.

(iv) If Y has only one complement, then $Y = \{0\}$ or $Y = X$.

(v) If $\operatorname{codim} Y > \dim Z$, then $Y + Z \neq X$.

(vi) If $\dim Y > \operatorname{codim} Z$, then $Y \cap Z \neq \{0\}$.

4.20 Let $(d_1, \ldots, d_r)$ be a sequence of natural numbers, and let X be an n-dimensional vector space. There exists a direct sum decomposition

$$X = Y_1 \dotplus \cdots \dotplus Y_r$$

with $\dim Y_j = d_j$, all j, if and only if $\sum_j d_j = n$.

4.21 Let d be any scalar-valued map, defined on the collection of all linear subspaces of a finite-dimensional vector space X, that satisfies the following two conditions: (i) $Y \cap Z = \{0\} \implies d(Y + Z) = d(Y) + d(Z)$; (ii) $\dim Y = 1 \implies d(Y) = 1$.

Prove that $d(Y) = \dim(Y)$ for every linear subspace of X.

4.22 Prove that the cartesian product $Y_1 \times \cdots \times Y_r$ of vector spaces, all over the same scalar field $\mathbb{F}$, becomes a vector space under *pointwise* or **slotwise** addition and multiplication by a scalar.

This vector space is called the **product space** with **factor**s $Y_1, \ldots, Y_r$.

Elimination in vector spaces

In the discussion of the (4.5)Basis Selection Algorithm, we left unanswered the unspoken question of just how one would tell which columns of $W \in L(\mathbb{F}^m, X)$ are bound, hence end up in the resulting 1-1 map V.

The answer is immediate in case $X \subset \mathbb{F}^r$ for some r, for then W is just an $r \times m$-matrix, and elimination does the trick since it is designed to determine the bound columns of a matrix. It works just as well when X is, more generally, a subset of $\mathbb{F}^T$ for some set T, as long as T is finite, since we can then apply elimination to the 'matrix'

$$W = (w_j(t) : (t, j) \in T \times \underline{m}) \tag{4.29}$$

whose rows are indexed by the (finitely many) elements of T.

Elimination even works when T is not finite, since looking for a pivot row in the matrix (4.29) with *infinitely* many rows is only a *practical* difficulty. If τ_i is the row 'index' of the pivot row for the ith bound column of W, $i = 1{:}r$, then we know that W has the same nullspace as the (finite-rowed) matrix $(w_j(\tau_i) : i = 1{:}r, j = 1{:}m)$. This proves, for arbitrary T, the following important

(4.30) Proposition: For any $W \in L(\mathbb{F}^m, \mathbb{F}^T)$, there exists a sequence $(\tau_1, \ldots, \tau_r)$ in T, with r equal to the number of bound columns in W, so that $\operatorname{null} W$ is equal to the nullspace of the matrix $(w_j(\tau_i) : i = 1{:}r, j = 1{:}m)$.

In particular, W is 1-1 if and only if the matrix $(w_j(\tau_i) : i, j = 1{:}m)$ is invertible for some sequence $(\tau_1, \ldots, \tau_m)$ in T.

If T is not finite, then we may not be able to determine in finite time whether or not a given column is bound since we may have to look at infinitely many rows not yet used as pivot rows. The only efficient way around this is to have W given to us in the form

$$W = UA,$$

with U some 1-1 column map, hence A a matrix. Under these circumstances, the kth column of W is free if and only if the kth column of A is free, and the latter we can determine by elimination applied to A.

Indeed, if U is 1-1, then both W and A have the same nullspace, hence, by (4.3)Lemma, the kth column of W is bound if and only if the kth column of A is bound.

As an example, consider $W = [w_1, w_2, w_3, w_4]$, with $w_j : \mathbb{R} \to \mathbb{R} : t \mapsto \sin(t-j)$, $j = 1,2,3,4$. Hence, by the addition formula,

$$W = UA, \quad \text{with } U := [\sin, \cos], \quad A := \begin{bmatrix} \cos(-1) & \cos(-2) & \cos(-3) & \cos(-4) \\ \sin(-1) & \sin(-2) & \sin(-3) & \sin(-4) \end{bmatrix},$$

and we see at once that U is 1-1 (e.g. from the fact that $QU = \mathrm{id}_2$, with $Q : f \mapsto (f(\pi/2), f(0))$). We also see at once that the first two columns of A are bound (e.g., since $\cos(1)\cos(2) < 0$ while $\sin(1)\sin(2) > 0$), hence the remaining columns of A must be free (since there are no rows left to bind them). Consequently, the first two columns of W are bound, while the last two columns are free.

Note that, necessarily, U is a basis for $\operatorname{ran} W$ since $W = UA$ implies that $\operatorname{ran} W \subset \operatorname{ran} U$, hence having two columns of W bound implies that $2 \le \dim \operatorname{ran} W \le \dim \operatorname{ran} U \le \#U = 2$, and so U is 1-1 onto $\operatorname{ran} W$.

In general, it may be hard to find such a handy factorization $W = UA$ for given $W \in L(\mathbb{F}^m, X)$. In that case, we may have to *discretize* our problem by finding somehow some $Q \in L(X, \mathbb{F}^n)$ that is 1-1 on $\operatorname{ran} W$. With such a 'data map' Q in hand, we know that $\operatorname{null} W$ equals the nullspace of the *matrix* QW. In particular, the kth column of W is bound if and only if the kth column of the *matrix* QW is bound, and elimination applied to QW will ferret out all those columns.

The need for suitable 'data maps' here in the general case is one of many reasons why we now turn to the study of this second way of connecting our vector space X to some coordinate space, namely via linear maps from X to $\mathbb{F}^n$.

4.23 For each of the following column maps $V = [v_1, \ldots, v_r]$ into the vector space Π_4 of all real polynomials of degree ≤ 4, determine whether or not it is 1-1 and/or onto.

(a) $[()^3 - ()^1 + 1, ()^2 + 2()^1 + 1, ()^1 - 1]$; (b) $[()^4 - ()^1, ()^3 + 2, ()^2 + ()^1 - 1, ()^1 + 1]$; (c) $[1 + ()^4, ()^4 + ()^3, ()^3 + ()^2, ()^2 + ()^1, ()^1 + 1]$.

4.24 For each of the specific column maps $V = [f_j : j = 0{:}r]$ given below (with f_j certain real-valued functions on the real line), determine which columns are bound and which are free. Use this information to determine (i) a basis for $\operatorname{ran} V$; and (ii) the smallest n so that $f_n \in \operatorname{ran}[f_0, f_1, \ldots, f_{n-1}]$.

(a) $r = 6$, and $f_j : t \mapsto (t-j)^2$, all j.

(b) $r = 4$ and $f_j : t \mapsto \sin(t-j)$, all j.

(c) $r = 4$ and $f_j : t \mapsto \exp(t-j)$, all j. (If you know enough about the exponential function, then you need not carry out any calculation on this problem.)

4.25 Assume that $\tau_1 < \cdots < \tau_{2k+1}$. Prove that $W = [w_0, \ldots, w_k]$ with $w_j : t \mapsto (t - \tau_{j+1}) \cdots (t - \tau_{j+k})$ is a basis for Π_k. (Hint: Consider QW with $Q : p \mapsto (p(\tau_{k+1+i}) : i = 0{:}k)$.)

4.26 Assume that $(\tau_1, \ldots, \tau_{2k+1})$ is nondecreasing. Prove that $W = [w_0, \ldots, w_k]$ with $w_j : t \mapsto (t - \tau_{j+1}) \cdots (t - \tau_{j+k})$ is a basis for Π_k if and only if $\tau_k < \tau_{k+1}$.

4.27 T/F

(a) If one of the columns of a column map is 0, then the map cannot be 1-1.

(b) If the column map V into $\mathbb{R}^n$ is 1-1, then V has at most n columns.

(c) If the column map V into $\mathbb{R}^n$ is onto, then V has at most n columns.

(d) If a column map fails to be 1-1, then it has a zero column.

5. The inverse of a basis, and interpolation

Data maps

There are two ways to connect a given vector space X with the coordinate space $\mathbb{F}^n$ in a linear way, namely by a linear map from $\mathbb{F}^n$ to X, and by a linear map to $\mathbb{F}^n$ from X. By now, we are thoroughly familiar with the first kind, the column maps. It is time to learn something about the other kind.

A very important example is the inverse of a basis $V : \mathbb{F}^n \to X$ for the vector space X, also known as the **coordinate map** for that basis because it provides, for each $x \in X$, its **coordinates with respect to the basis**, i.e., the n-vector $a := V^{-1}x$ for which $x = Va$. In effect, every *invertible* linear map from X to $\mathbb{F}^n$ is a coordinate map, namely the coordinate map for its inverse. However, (nearly) every linear map from X to $\mathbb{F}^n$, invertible or not, is of interest, as a means of extracting numerical information from the elements of X. For, we can, offhand, only compute with numbers, hence can 'compute' with elements of an abstract vector space only in terms of numerical data about them.

Any linear map from the vector space X to $\mathbb{F}^n$ is necessarily of the form

$$f : X \to \mathbb{F}^n : x \mapsto (f_i(x) : i = 1{:}n),$$

with each $f_i = {e_i}^{\mathrm{t}} \circ f$ a **linear functional** on X, i.e., a scalar-valued linear map on X. Here are some standard examples.

5.1 For each of the following maps, determine whether or not it is a linear functional. (a) $\Pi_k \to \mathbb{R} : p \mapsto \deg p$; (b) $\mathbb{R}^3 \to \mathbb{R} : x \mapsto 3x_1 - 2x_3$; (c) $C[a\mathinner{\ldotp\ldotp}b] \to \mathbb{R} : f \mapsto \max_{a \le t \le b} f(t)$; (d) $C[a\mathinner{\ldotp\ldotp}b] \to \mathbb{R} : f \mapsto \int_a^b f(s)w(s)\,\mathrm{d}s$, with $w \in C[a\mathinner{\ldotp\ldotp}b]$; (e) $C^{(2)}(\mathbb{R}) \to \mathbb{R} : f \mapsto a(t)D^2 f(t) + b(t)Df(t) + c(t)f(t)$, for some functions a, b, c defined on $[a\mathinner{\ldotp\ldotp}b]$ and some $t \in [a\mathinner{\ldotp\ldotp}b]$. (f) $C^{(2)}(\mathbb{R}) \to C(\mathbb{R}) : f \mapsto aD^2 f + bDf + cf$, for some $a, b, c \in C(\mathbb{R})$.

Assume that X is a space of functions, hence X is a linear subspace of $\mathbb{F}^T$ for some set T. Then, for each $t \in T$,

$$\delta_t : X \to \mathbb{F} : x \mapsto x(t)$$

is a linear functional on X, the linear functional of evaluation at t. For any n-sequence $s = (s_1, \ldots, s_n)$ in T,

$$X \to \mathbb{F}^n : f \mapsto (f(s_1), \ldots, f(s_n))$$

is a standard linear map from X to $\mathbb{F}^n$.

If, more concretely, X is a linear subspace of $C^{(n-1)}[a\mathinner{\ldotp\ldotp}b]$ and $s \in [a\mathinner{\ldotp\ldotp}b]$, then

$$X \to \mathbb{F}^n : f \mapsto (f(s), Df(s), \ldots, D^{n-1}f(s))$$

is another standard linear map from such X to $\mathbb{F}^n$.

Finally, if $X = \mathbb{F}^m$, then any linear map from X to $\mathbb{F}^n$ is necessarily a matrix. But it is convenient to write this matrix in the form A^{t} for some $A \in \mathbb{F}^{n \times m}$, as such A^{t} acts on X via the rule

$$X \mapsto \mathbb{F}^n : x \mapsto A^{\mathrm{t}}x = (A(:, j)^{\mathrm{t}}x : j = 1{:}n).$$

Because of this last example, we will call all linear maps from a vector space to a coordinate space **row maps**, and use the notation

$$\Lambda^{\mathrm{t}} : X \to \mathbb{F}^n : x \mapsto (\lambda_i x : i = 1{:}n) =: [\lambda_1, \ldots, \lambda_n]^{\mathrm{t}}x, \tag{5.1}$$

calling the linear functional λ_i the ith **row** of this map. We will also call such maps **data maps** since they extract numerical information from the elements of X. There is no hope of doing any practical work with the vector space X unless we have a ready supply of such data maps on X. For, by and large, we can only compute with numbers.

(5.2)Proposition: If $\Lambda^{\mathrm{t}} = [\lambda_1, \lambda_2, \ldots, \lambda_n]^{\mathrm{t}} : X \to \mathbb{F}^n$ and $B \in L(U, X)$, then $\Lambda^{\mathrm{t}}B = [\lambda_1 B, \ldots, \lambda_n B]^{\mathrm{t}}$.

