Part III: Higher Dimensions.

21. The Fundamental Dichotomy. A square matrix A of any size (say, $n \times n$) is *invertible* if a matrix A^{-1} can be found such that $AA^{-1} = A^{-1}A = I_n$ is the $n \times n$ identity matrix. In that case, any equation of the form $AX = C$ (with C being a given column) has exactly one solution, namely $X = A^{-1}C$. In particular, the nullspace $\mathcal{N}(A)$ contains only zero.

The friendliest square matrices are the elementary ones: each of them is obtained from the identity matrix by a one-step row operation (remember?) and has an inverse that can be instantly written down. Any product of elementary matrices is explicitly invertible.

Much of linear algebra hinges on a surprising connection between two seemingly unrelated problems:

(i) Find an n-column $X \neq 0$ such that $AX = 0$.
(ii) Factor A into a product of elementary matrices.

Obviously (ii) excludes (i) and vice versa, because invertible matrices cannot have non-trivial nullspaces.

Theorem 1: For any square matrix A, one and only one of these problems is solvable.

Now this is a rather sweeping statement considering that n might be 100^{100} or more. What if tomorrow's technology finds a counter-example? Well, let us consider that possibility. *If* the theorem does not always hold, then surely there will be a *smallest* n for which it fails. Such a minimal counter-example would be an $n \times n$-matrix A with $\mathcal{N}(A) = 0$, refusing to be factored into elementaries. On the other hand, matrices with fewer than n rows would still obey the theorem.

To start with, let us check $n = 1$. Then $\mathcal{N}(A) = 0$ would force A to be a non-zero 1×1-matrix, i.e. an elementary matrix of type D. So far, so good.

Therefore take $n > 1$, and let A be an $n \times n$-matrix such that $\mathcal{N}(A) = 0$. Then the first column of A must be non-zero or else $(1, 0, \ldots, 0)^T$ would be in the nullspace of A. Hence we can row-reduce the equation $AX = 0$ to the equivalent $MAX = 0$, where $M = E_r \cdots E_1$ is a product of elementaries and the matrix MA has $(1, 0, \ldots, 0)^T$ as its first column. Explicitly,

$$MAX = \begin{bmatrix} 1 & R' \\ 0 & A' \end{bmatrix} \begin{bmatrix} x \\ X' \end{bmatrix} = \begin{bmatrix} x + R' \bullet X' \\ A'X' \end{bmatrix},$$

where R' is an n'-row, X' and 0 are n'-columns, x is a scalar, and A' is an $n' \times n'$-matrix, with $n' = n - 1$. This smaller matrix A' must fall into one of two patterns:

(i) There is an $X' \neq 0$ with $A'X' = 0$. Then we could get $MAX = 0$ with a non-zero X made up from X' and $x = -R \bullet X'$. But $\mathcal{N}(MA) = \mathcal{N}(A) = 0$, hence this case is impossible.

(ii) A' is a product of elementary matrices. Then we can row-reduce MA further (working only on the last n' rows) until the block occupied by A' is just $I_{n'}$. Finally we would use the 1's on the diagonal to reduce R' to zero as well, and obtain $E_s \cdots E_1 A = I_n$, as advertised.

Rectangular matrices. We now abandon our policy of sticking to square matrices and allow A to be an $m \times n$-matrix, that is, having m rows and n columns with m and n not necessarily equal. This does not make the discussion more difficult but simply more awkward, especially with regard to multiplication. If B is an $p \times m$-matrix, the product BA, defined by dotting rows of B with columns of A, will of course be a $p \times n$-matrix.

Most questions concerning rectangular matrices can be answered by what we know about square ones. For example, if A has fewer rows than columns (i.e. $m < n$), we can imagine it enlarged to an $n \times n$-matrix A^+ by additional rows of zeroes. Then obviously $\mathcal{N}(A) = \mathcal{N}(A^+) \neq 0$, because with its zero rows A^+ cannot possibly be invertible.

Corollary: If A has more columns than rows, then $\mathcal{N}(A) \neq 0$.

22. Full Reduction. In this lesson we shall use the expression "elementary product" to refer to a product of elementary matrices. By Theorem 1, we could just as well say "invertible matrix" or "non-singular square matrix". However, we do not want to invoke that theorem, since the following method yields, among other things, an independent derivation of it.

As before, the technique of *elimination* consists in using a sequence of elementary row operations (= left multiplication by an elementary product) to simplify a given matrix A for certain purposes. We shall here describe a particularly thorough version of it, known as *full reduction*, which looks at the columns of A one at a time and tries to turn as many as possible into columns of I_m, as follows.

If $C = (c_1, \ldots, c_m)^T \neq 0$ is an m-column, pick an index k such that $c_k \neq 0$. Now apply $D_k(c_k^{-1})$ to turn the k-th coordinate into 1, and then reduce the other coordinates to 0 by elementary operations of type $G_{ik}(\cdot)$. This sequence of operations is known as "pivoting C by the k-th row". It amounts to left multiplication by an elementary product M with the property that (a) $MC = I_m^{(k)}$ is the k-th column of I_m, and (b) $MK = K$ for any column K whose k-th coordinate is 0.

The matrix A is said to be *fully reduced* if the following two numbers are equal:

$\nu(A)$=the number of distinct columns of I_m contained in A;
$\mu(A)$=the number of non-zero rows of A.

Note: We always have $\nu(A) \leq \mu(A)$, since every column of I_m, with its 1, causes a different row to be non-zero. Full reduction simply means that there are no additional non-zero rows. Thus, if we were allowed some reordering of rows *and columns*, a fully reduced $A' = MA$ would appear in block form as

$$A' = \begin{bmatrix} I_r & F \\ 0 & 0 \end{bmatrix} \tag{1}$$

where $r = \nu(A) = \mu(A)$. Without actually reshuffling the columns $C'_1, \ldots, C'_n$ of A', this amounts to partitioning them into two sets: the "bound" columns $C'_{\beta(1)}, \ldots, C'_{\beta(r)}$, which look like the first r columns of I_m, and the "free" columns $C'_{\varphi(1)}, \ldots, C'_{\varphi(n-r)}$, which are unspecified except that they must be zero below the r-th row. Free columns exist whenever $r < n$, and only then.

Lemma : Given any A, we can find an elementary product M, such that $A' = MA$ is fully reduced.

This results immediately from the important observation that, *if $\nu(A) < \mu(A)$, there is an elementary product M such that $\nu(MA) > \nu(A)$.*

Indeed, by suitably ordering the rows, we may assume that A contains the first $\nu(A)$ columns of I_m. Then, if $\nu(A) < \mu(A)$, there must be a non-zero k-th row with $k > \nu(A)$. If $a_{kl} \neq 0$, we can pivot the l-th column by the k-th row to change it into $I_m^{(k)}$. This does not affect the previously existing $I_m^{(i)}$, for $i = 1, \ldots, \nu(A)$, because their k-th coordinates are 0. Thus we obtain a net increase in the number of such columns, as claimed.

Finally we are ready to investigate the equation $AX = 0$, which is equivalent to $MAX = 0$ for any invertible M. If M equals the elementary product provided by the lemma, this equation just says

$$x_1 C'_1 + \cdots + x_n C'_n = 0, \tag{2}$$

where C'_k stands for the k-th column of A'. Now suppose that $C'_{\varphi(j)} = (z_{1j}, \ldots, z_{rj}, 0, \ldots, 0)^T$ is a *free* column of A'. Then the obvious equality

$$C'_{\varphi(j)} = z_{1j} C'_{\beta(1)} + \cdots + z_{rj} C'_{\beta(r)} \tag{3}$$

gives a solution of (2) with $x_{\varphi(j)} = 1$. This happens for every particular free index $\varphi(j)$, for $j = 1, \ldots, n-r$. If A is square ($m = n$), we either have such solutions ($r < n$), or we have $r = m = n$ and hence $MA = I_m$, in which case A itself is an elementary product.

23. Dimension of Subspaces. In lesson 13, we classified matrices according to the dimension of their nullspaces, which appeared geometrically as lines or planes. To carry this over to larger matrices, and still be sure of what we mean, we have to free the relevant concepts of their visual underpinning. If we want to keep some of the geometric language, we have to anchor it carefully in the formal, non-visual world of algebra.

In the following, the word "vector" will mean an n-tuple $(a_1, \ldots, a_n)$ of numbers. The set of all vectors will be denoted by $\mathbf{R}^n$. The integer n will be kept fixed. A non-empty subset $\mathcal{V} \subseteq \mathbf{R}^n$ is called a *subspace*, if $V, W \in \mathcal{V}$ implies $aV + bW \in \mathcal{V}$, for any scalars a, b.

At this point you should pause and try to see why the subspaces of $\mathbf{R}^3$ (apart from the trivial $\{0\}$ and the whole space) are precisely the ones which, pictured as sets of coordinatized points, show up as lines or planes.

It is easy to check that the nullspace of any $m \times n$-matrix is a subspace. Another interesting type of subspace is the set of all *linear combinations* $c_1V_1 + \cdots + c_rV_r$ of some given r-tuple of vectors. It is called the *span* of $V_1, \ldots, V_r$, and is obviously the smallest subspace containing these vectors. The set $\{V_1, \ldots, V_r\}$ itself is said to be *independent*, if no non-trivial linear combination of it is zero. For $r > 1$, this means that none of these vectors is a linear combination of the other $(r-1)$.

The most important relation between these concepts is:

(1) *An independent subset of the span of r vectors cannot have more than r elements.*

How come? Suppose $W_1, \ldots, W_s$ are in the span of $V_1, \ldots, V_r$; say $W_j = a_{1j}V_1 + \cdots + a_{rj}V_r$, for $j = 1, \ldots, s$. Consider the linear combination

$$x_1W_1 + \cdots + x_sW_s = (a_{11}x_1 + \cdots + a_{1s}x_s)V_1 + \cdots + (a_{r1}x_1 + \cdots + a_{rs}x_s)V_s.$$

If $s > r$, the matrix (a_{ij}) involved here has more columns than rows, and hence the corollary at the end of lesson 21 guarantees the existence of a non-trivial s-tuple $x_1, \ldots, x_s$ such that all this is zero.

For the rest of this page, let $\mathcal{V} \neq \{0\}$ be a fixed subspace of $\mathbf{R}^n$. An independent subset $\{W_1, \ldots W_k\} \subset \mathcal{V}$ is called a *basis* of $\mathcal{V}$, if it spans all of $\mathcal{V}$, i.e. if every $W \in \mathcal{V}$ can be written in the form $W = c_1W_1 + \cdots + c_kW_k$. Because of the independence it can never happen that $c_1W_1 + \cdots + c_kW_k = c_1'W_1 + \cdots + c_k'W_k$ unless $c_i = c_i'$ throughout (otherwise subtract one expression from the other and get a dependence relation). Hence the k-tuple $(c_1, \ldots, c_k)$ is a unique label for W, and thus the basis $\{W_1, \ldots, W_k\}$ acts as a kind of "coordinate system" for the subspace $\mathcal{V}$.

You should pause to reflect that the standard coordinates of a vector $(a_1, \ldots, a_n) \in \mathbf{R}^n$ are obtained in just this way from the basis formed by all the rows of the identity matrix I_n. Can you also see, how any pair of non-parallel vectors in $\mathbf{R}^3$ forms a basis for some plane ?

By (1), no basis of $\mathcal{V}$ can have more elements than any other, i.e. the number of vectors in any basis is the same; it is called the *dimension* of $\mathcal{V}$, written $\dim \mathcal{V}$. We now need to convince ourselves that bases are always available.

(2) *Any independent subset $\{W_1, \cdots, W_s\}$ of $\mathcal{V}$ is contained in a basis of $\mathcal{V}$.*

In fact, if $\{W_1, \ldots, W_s\}$ is not already a basis, there must be a $W \in \mathcal{V}$ such that $\{W_1, \ldots, W_s, W\}$ is still independent (see?). So, keep adjoining more vectors $W_{s+1}, W_{s+2}, \ldots \in \mathcal{V}$ (if you can), while maintaining the independence of your collection. By (1), this process cannot go beyond a total of n vectors, because $\mathbf{R}^n$ itself is the span of the rows of I_n. At some point, therefore, your set $\{W_1, \ldots, W_{s+p}\}$ stops being enlargeable and hence must be a basis.

This result also shows that $\dim \mathcal{V}$ is a meaningful measure of the "size" of $\mathcal{V}$. More precisely, if $\mathcal{V}$ contains a smaller subspace $\mathcal{V}'$, we can enlarge a basis of $\mathcal{V}'$ to one of $\mathcal{V}$, thus proving that $\dim \mathcal{V}' < \dim \mathcal{V}$.

24. Bases by Elimination. With any $m \times n$-matrix A, we now associate two subspaces $\mathcal{N}(A)$ and $\mathcal{R}(A)$ of $\mathbf{R}^n$ and one subspace $\mathcal{C}(A)$ of $\mathbf{R}^m$. Of course, $\mathcal{N}(A)$ is the nullspace. The *row-space* $\mathcal{R}(A)$ and the *column-space* $\mathcal{C}(A)$ are just the spans of the rows and the columns of A, respectively.

We shall see that these spaces play important roles in problems concerning a given A. Conversely, whenever we face a problem about subspaces and dimension, our usual response will be to create a suitable $m \times n$-matrix A such that the subspace in question turns up either as $\mathcal{R}(A)$ or $\mathcal{C}(A)$, or even as $\mathcal{N}(A)$. Then we apply elimination (using elementary row operations) until we arrive at a matrix $A' = MA$ from which the desired information can be read off.

In doing this, we need to keep in mind the following basic relations between A and $A' = MA$ for invertible M:

$$\mathcal{N}(A') = \mathcal{N}(A), \qquad \mathcal{R}(A') = \mathcal{R}(A), \qquad \dim \mathcal{C}(A') = \dim \mathcal{C}(A). \tag{1}$$

The first and last of these are two faces of the same coin, namely $x_1C_1 + \cdots x_nC_n = 0 \iff x_1MC_1 + \cdots + x_nMC_n = 0$, where $C_1, \ldots C_n$ are the columns of A. This shows that left multiplication with M neither creates nor destroys any linear relations between individual columns. In particular, *$\mathcal{B}$ is a basis of $\mathcal{C}(A)$ if and only if $M\mathcal{B}$ is a basis of $\mathcal{C}(A')$.*

For the middle part of (1), we note that every row of MA is a linear combination of the rows $R_1, \ldots, R_m$ of A (with coefficients from the corresponding row of M), and hence $\mathcal{R}(MA)$ is contained in $\mathcal{R}(A)$. This is just matrix multiplication and works for any M and any A. If M is invertible (as it is here), we also get the opposite inclusion, because $A = M^{-1}(MA)$.

If we have a fully reduced $A' = MA$ (cf. lesson 22), it is now easy to find bases for the subspaces associated with A. The non-zero rows of A' are clearly independent (why?) and hence form a basis of $\mathcal{R}(A') = \mathcal{R}(A)$. Just as clearly, the bound columns $C'_{\beta(j)} = MC_{\beta(j)}$ of A' form a basis of $\mathcal{C}(A')$; hence the corresponding columns $C_{\beta(1)}, \ldots, C_{\beta(r)}$ of the original matrix form a basis of $\mathcal{C}(A)$. Note in particular that both $\mathcal{R}(A)$ and $\mathcal{C}(A)$ have the same dimension r.

Finally, the linear relations (3) of lesson 22 yield a set of $n-r$ elements of $\mathcal{N}(A)$ which are independent, because at every free index $\varphi(j)$ exactly one of them is non-zero (see?). They do form a basis of $\mathcal{N}(A)$, because $n - r$ is the dimension of this space, according the following theorem.

Theorem 2:

$$\dim \mathcal{R}(A) = \dim \mathcal{C}(A) = n - \dim \mathcal{N}(A). \tag{2}$$

The first of these equalities has just been explained; it is restated here only for emphasis. For the second one, we start with a basis $\{W_1, \ldots, W_k\}$ of the nullspace and extend it to a basis $\{W_1, \ldots, W_n\}$ of $\mathbf{R}^n$. If we can show that $\{AW_{k+1}, \ldots, AW_n\}$ is a basis of $\mathcal{C}(A)$, we will have proved our point. Now, any dependence relation $0 = c_{k+1}AW_{k+1} + \cdots + c_nAW_n$ would mean that $c_{k+1}W_{k+1} + \cdots + c_nW_n \in \mathcal{N}(A)$, and hence imply a dependence relation among the W_i (see?). Moreover, since every column of I_n is a linear combination of the W_i, every column of A is a linear combination of the AW_i. But the first k of these are 0, and hence the remaining ones span $\mathcal{C}(A)$.

Linear equations. Given an m-column C, it is often required to find all X satifying the equation

$$AX = C. \tag{3}$$

Clearly such X exist if and only if $C \in \mathcal{C}(A)$, in which case the equation is said to be *consistent*. Suppose we have located one such solution $X = V$; then any other solution must be of the form $V + W$ with $W \in \mathcal{N}(A)$ (see?). Hence if $W_1, \ldots, W_s$ form a a basis of $\mathcal{N}(A)$, the *general* solution of (3) is

$$X = V + t_1W_1 + \cdots + t_sW_s, \tag{4}$$

with s free parameters $t_1, \ldots, t_s$.

25. Orthogonality Again. You are invited to reconsider lesson 20, but this time imagine that A is an $m \times 3$-matrix, with $m > 3$, and to ask yourself what changes (if any) would be required in the text.

A fair bit of rethinking is certainly necessary. In addition to certain 3-vectors and 3×3-matrices, which stay in the game, the lesson would now also refer to m-vectors as well as $m \times 3$ and $3 \times m$-matrices. (You should spend some time to figure out which is which.) However, you will find that in the actual *wording* no change is required.

You might say that the reference to Pythagoras is wrong, because there is no geometry beyond $3D$. To this we reply that we did not need *geometry* as such in formula (7), but only the obvious algebraic fact that

$$V \bullet W = 0 \qquad \Longrightarrow \qquad (V+W) \bullet (V+W) = V \bullet V + W \bullet W. \tag{1}$$

Ascribing that to Pythagoras may be a historical mistake but certainly not a logical one. To stay out of trouble we refrain from thinking of $|V| = \sqrt{V \bullet V}$ as a length; if a name is needed we call this number the *norm* of V.

Thus the argument for the optimality of the "least squares" solution remains valid and so does its conclusion: if X is a 3-column such that $A^TAX = A^TC$ (this X is not hard to find once we have $A = GS$) then $|AX - C|$ is minimal. Typical applications of this technique tend to have m much bigger than 3 — cf. "linear regression".

Of course, there is no reason to limit the number of columns to 3. If $A = \big[V_1, \ldots, V_n\big]$ is an $m \times n$-matrix, the Gram-Schmidt process creates a unique column-rectified counterpart $G = \big[W_1, \ldots, W_n\big]$ such that $A = GS$, where S is upper triangular and unipotent. After $W_1 = V_1$, its $(k+1)$-st step is as follows:

$$W_{k+1} = V_{k+1} - \mathrm{proj}_{W_1}(V_{k+1}) - \cdots - \mathrm{proj}_{W_k}(V_{k+1}). \tag{2}$$

As in lesson 8, $\mathrm{proj}_W(V) = tW$ is defined by the condition $(V - tW) \bullet W = 0$, which here should be regarded as purely algebraic.

One aspect that was not stressed in lesson 20 is this: $GY = 0 \iff DY = 0$, and this happens with $Y \neq 0$ (i.e. A and G are singular) if and only if one of the W_i is zero. In this way Gram-Schmidt also monitors the non-singularity of A.

Lengths and angles. Let $\mathcal{V}$ be a 2-dimensional subspace of $\mathbf{R}^m$, say, the span of two vectors V_1 and V_2. By running this basis of $\mathcal{V}$ through the GS-machine if necessary, we can replace it by an orthogonal pair W_1, W_2, and by possibly rescaling these, we may assume that $W_i \bullet W_i = 1$. Then clearly

$$(x_1W_1 + x_2W_2) \bullet (y_1W_1 + y_2W_2) = x_1y_1 + x_2y_2, \tag{3}$$

for any pair of vectors $X = x_1W_1 + x_2W_2$ and $Y = y_1W_1 + y_2W_2$ in $\mathcal{V}$ (see?). The right hand side of this formula should look familiar: it is the dot-product of two ordinary (geometric) vectors $\hat{X} = (x_1, x_2)$ and $\hat{Y} = (y_1, y_2)$ in the ordinary Euclidean plane $\mathbf{R}^2$. If ϕ denotes the angle between these vectors, we have the amazing formula

$$X \bullet Y = |\hat{X}| \cdot |\hat{Y}| \cdot \cos\phi, \tag{4}$$

amazing because the left hand side denotes a formal algebraic operation in a socalled "space" of many socalled "dimensions", while the right is firmly anchored in good old trigonometry.

In effect we have made a precise, mathematical *picture* of the subspace $\mathcal{V}$, even though its environment $\mathbf{R}^m$ as a whole remains invisible. Obviously we can create an analogous image for any 3-dimensional subspace. Measurements of lengths and angles taken in different images cannot contradict one another, because they are regulated by the dot-product in $\mathbf{R}^m$. Thus there is no possibility of logical or numerical inconsistency, if we refer to the the square root $|X|$ of $X \bullet X$ as the *length* of X, or to

$$\phi = \cos^{-1} \frac{X \bullet Y}{|X| \cdot |Y|}, \tag{5}$$

as the *angle* between X and Y.

26. The Spectral Theorem. Back to square matrices and the eigenvector hunt. Our previous findings suggest that any *symmetric* $n \times n$-matrix A should have n mutually orthogonal eigenvectors; that is, it should be possible to find an orthogonal matrix Q, such that $Q^T AQ$ is diagonal. In this lesson we shall see that this is indeed the case.

Theorem 3: For every (real) symmetric A there is an orthogonal Q such that $Q^T AQ = D$ is diagonal.

But instead of taking the former detour via determinants and characteristic polynomials (which is impractical for $n > 3$), we shall directly attack and decimate the off-diagonal entries by aptly chosen orthogonal similarities. Let A and Q be the second and third matrices (respectively) on the left hand side of the formula

$$Q^T AQ = \begin{bmatrix} M^T & 0 \\ 0 & I_{n-2} \end{bmatrix} \begin{bmatrix} A' & X \\ X^T & A'' \end{bmatrix} \begin{bmatrix} M & 0 \\ 0 & I_{n-2} \end{bmatrix} = \begin{bmatrix} M^T A'M & M^T X \\ X^T M & A'' \end{bmatrix} = B \tag{1}$$

in which M is an orthogonal 2×2-matrix, and A' is the upper left 2×2-submatrix of A. Since A' is symmetric, we can choose M in such a way that $M^T A'M = D'$ is diagonal; that is, in $B = Q^T AQ$ we have $b_{12} = b_{21} = 0$.

The problem is that we have killed only one off-diagonal pair, and that the newly created zeroes might be messed up again as we try to continue the process. A closer look, however, reveals that there *is* a net gain: the total off-diagonal mass is down and will stay down.

To see this, let $\sigma(A)$ denote the sum of the squares of off-diagonal entries. Then, with reference to the blocks shown in (1), we clearly have

$$\sigma(A) = \sigma(A') + 2|X|^2 + \sigma(A''), \tag{2}$$

where $|X|^2$ stands for the sum of the squares in the $2 \times (n-2)$-block X. When we look at $\sigma(B)$ in the same way, we find $\sigma(D') = \sigma(A') - 2a_{12}^2$ in the first term and and an unchanged last term $\sigma(A'')$. The middle term has apparently changed to $|M^T X|^2$ — but that still equals $|X|^2$, because the orthogonal M^T preserves the length of every column of X (see?). Altogether: $\sigma(B) = \sigma(A) - a_{12}^2$.

Of course, a similar move can be made with any off-diagonal index pair i, j in the place of $1, 2$. To fit this literally into the scheme described above, we would first change A to $P^T AP$ with a permutation matrix P which swaps the indices i and 1 as well as j and 2 (if necessary). We state our findings formally.

Lemma : For every symmetric A we can find an orthogonal Q such that $\sigma(Q^T AQ) = \sigma(A) - a_{ij}^2$, where i, j is any given index pair with $i \neq j$.

To prove the theorem, consider the function $f_A(Q) = \sigma(Q^T AQ)$ as Q ranges over all possible orthogonal $n \times n$-matrices. Let v_0 be the minimal value of this function, and take Q_0 such that $f_A(Q_0) = v_0$. Now, if it were the case that $v_0 > 0$, then some off-diagonal element would have to be $\neq 0$, and the lemma would yield a Q_1 with $f_A(Q_1) < v_0$ — an impossibility. Hence all off-diagonal elements in $Q_0^T AQ_0$ must be dead.

Unfortunately this nifty argument has a gap: how do we know that there is a Q_0 which minimizes f_A? There are many decent functions, for instance $f(x) = \exp(-x^2)$, which never attain a minimum. But we are lucky in that the set of orthogonal $n \times n$-matrices is closed and bounded in $\mathbf{R}^{n^2}$. Since every column of such a matrix Q must have length 1, the matrices themselves (regarded as n^2-tuples) lie on a sphere of radius $\sqrt{n}$: the set is *bounded*. Furthermore, the equation $Q^T Q = I$ clearly survives any limit operation; that is, a limit of orthogonal matrices is still orthogonal: the set is *closed*.

One of the most basic lemmas of real analysis says that a continuous real-valued function on a closed and bounded subset of $\mathbf{R}^N$ always attains a minimum — and our argument is saved.

You should now look back at lesson 17 and check that its second half carries over without change.

27. The Minimal Polynomial. Where have all the lambdas gone? No, we have not lost interest in eigenvalues, but we are lacking one of the basic tools used in low dimensions: the characteristic polynomial $p_A(\lambda) = \det(\lambda I - A)$. So far, we do not even have determinants. Not that these items are unavailable; they are just excessively cumbersome.

It would take us too far afield to give accounts of the various iterative "direct" methods for hunting eigenvalues, and so we shall stick to polynomials, but use the *minimal* polynomial $m_A(\lambda)$ instead of the characteristic one. As defined in lesson 16,

$$m_A(\lambda) = c_0 + c_1\lambda + \cdots c_{s-1}\lambda^{s-1} + \lambda^s \tag{1}$$

is the smallest (in degree s) polynomial annihilating A, i.e. giving 0 when A is substituted for λ. Every eigenvalue λ_1 of A is clearly a root of any annihilating polynomial, hence of $m_A(\lambda)$. Conversely, if $m_A(\alpha) = 0$, then α is an eigenvalue. Indeed, $m_A(\lambda) = (\lambda - \alpha)h(\lambda)$ with $h(A) \neq 0$ by reason of degree; however, since $(A - \alpha I)h(A) = 0$ means $A\,h(A) = \alpha h(A)$, every non-zero column of $h(A)$ must be an eigenvector. Thus $m_A(\lambda)$ is as good for locating eigenvalues as the characteristic polynomial, which we are boycotting. The only problem with this approach is that we have (so far) no assurance that $s \leq n$.

We can move toward filling this gap at the same time as we start computing $m_A(\lambda)$. As in lesson 16, we pick a vector $V \neq 0$ and find a non-trivial dependence relation

$$x_0V^{(0)} + x_1V^{(1)} + \cdots + x_nV^{(n)} = 0, \tag{2}$$

where $V^{(k)} = A^kV$ for $k = 0, 1, \ldots, n$. In doing so, we make sure that the last non-zero term of our solution sequence $x_0, x_1, \ldots, x_n$ is $x_t = 1$ with $1 \leq t \leq n$ as small as possible. Then the polynomial $g(\lambda) = x_0 + x_1\lambda + \cdots + x_{t-1}\lambda^{t-1} + \lambda^t$ is of minimal degree with the property that $g(A)V = 0$. We refer to it as the *minimal polynomial of A acting on V*, and denote it by $m_{A|V}(\lambda)$. Obviously every $f(\lambda)$ such that $f(A)V = 0$, for instance $m_A(\lambda)$, is a multiple of $m_{A|V}(\lambda)$. This is the first step in the computation of $m_A(\lambda)$. We now pause to show that $s \leq n$, using $g(\lambda) = m_{A|V}(\lambda)$ in the following argument.

Putting $u = n - t$, we choose vectors $W_1, \ldots, W_u$ such that the matrix $\left[V^{(0)}, \ldots, V^{(t-1)}, W_1, \ldots, W_u\right] = M$ is invertible. By construction $AV^{(0)} = V^{(1)}$, etc., and $AV^{(t-1)} = -x_0V^{(0)} - \cdots - x_{t-1}V^{(t-1)}$. Therefore $AM = MB$, with

$$B = \begin{bmatrix} T & * \\ & \\ 0 & U \end{bmatrix} \qquad \text{and more generally} \qquad f(B) = \begin{bmatrix} f(T) & ** \\ & \\ 0 & f(U) \end{bmatrix}, \tag{3}$$

where $f(\lambda)$ is any polynomial, and T, U are certain matrices of size $t \times t$ and $u \times u$ respectively. The zero in the lower left is due to the fact that $AV^{(k)}$ does not involve any term in $W_1, \ldots, W_u$. Note that $g(T) = 0$, because $g(A)M$ starts with t zero columns. If $t = n$, we get $g(A) = 0$ — hooray! If $t < n$, the first t columns of $g(B)$ are zero; and so are the last u rows of $m_U(B)$ (because $m_U(U) = 0$). Therefore $g(B)m_U(B) = 0$, and the polynomial $f(\lambda) = g(\lambda)m_U(\lambda)$ annihilates A. On the other hand, since U is smaller than A, we may take it for granted that $m_U(\lambda)$ has degree $\leq u$. Thus $f(\lambda)$ has degree $\leq t + u = n$.

Now that we know it to be non-trivially solvable, we can subject the homogeneous linear equation

$$x_0I + x_1A + \cdots + x_nA^n = 0, \tag{4}$$

to the usual elimination procedure, again taking care that the solution sequence $x_0, x_1, \ldots, x_n$ has $x_r = 1$ as its last non-zero term, with r as small as possible. Of course, then $r = s$ and $x_0 + x_1\lambda + \cdots + \lambda^r$ is our minimal polynomial. Equation (4) is just the simultaneous enforcement of (2) for n independent choices $V = V_1, \ldots, V_n$ (the columns of I_n) stacked on top of one another, i.e. a system of n^2 equations in $n + 1$ unknowns. In practice, it is unwise to engage all these equations at the same time. Performing (2) with a single choice $V = V_1$ will usually yield $m_{A|V}(\lambda) = m_A(\lambda)$. If not, so much the better: we then have a divisor of lower degree and hence easier access to some of the eigenvalues. To get $m_A(\lambda)$ itself, we can always expand the system to include V_2, V_3, etc.

28. Diagonalization and Triangulation. It is time to pick up the thread of diagonalization outside of the special context of symmetric matrices. Suppose we have found the minimal polynomial $m_A(\lambda)$ and seen that it has a total of r distinct roots $\lambda_1, \ldots, \lambda_r$, what does that tell us about diagonalization?

Of course, the polynomial $l_A(\lambda) = (\lambda - \lambda_1) \cdots (\lambda - \lambda_r)$ is always a divisor of $m_A(\lambda)$. But if A is similar to a diagonal matrix D, the opposite division also works. Indeed, since $m_A(\lambda) = m_D(\lambda)$, the diagonal entries (eigenvalues) of D are just the scalars $\lambda_1, \ldots, \lambda_r$ each with a suitable multiplicity. Now

$$l_A(D) = (D - \lambda_1 I) \cdots (D - \lambda_r I) = 0, \tag{1}$$

because in every row of this product of diagonal matrices one of the factors is zero. Hence $l_A(\lambda)$ annihilates D and also A; being both a divisor and a multiple of $m_A(\lambda)$, it must coincide with it. (Incidentally, the characteristic polynomial of D would contain each factor $(\lambda - \lambda_k)$ not always just once, but as many times as λ_k occurs on the diagonal of D). As in lesson 14, we characterize diagonalizable matrices in terms of two conditions — the second of which looks quite different from its earlier counterpart (why?).

Theorem 4: An $n \times n$ matrix A is diagonalizable if and only if

(i*) *$m_A(\lambda)$ splits completely into a product of factors $(\lambda - \lambda_k)$ of degree one;*
(ii*) *each of these factors occurs only once, i.e. $m_A(\lambda) = (\lambda - \lambda_1) \cdots (\lambda - \lambda_s)$.*

By itself (i*) *means that A is similar to a triangular matrix.*

To establish these facts, we start with a slight variation of the paragraph around equation (3) in the preceding lesson, which we now modify as follows. Given $t < n$ independent eigenvectors $V_1, \ldots, V_t$ for the same eigenvalue λ_1, we concoct an invertible matrix $M = \big[V_1, \ldots, V_t, W_1, \ldots, W_u\big]$, and obtain a matrix $B = M^{-1}AM$ of the form displayed there, with the submatrices T and U in the upper left and lower right corners. Taking $g(\lambda) = (\lambda - \lambda_1)$, we then conclude as before that the polynomial $f(\lambda) = g(\lambda)m_U(\lambda)$ annihilates A. Hence $f(\lambda)$ is a multiple of $m_A(\lambda)$, which in turn is a multiple of $m_U(\lambda)$, because $m_A(U)$ is the $u \times u$ lower right submatrix of $m_A(B) = 0$. Therefore either

$$m_A(\lambda) = m_U(\lambda) \qquad \text{or} \qquad m_A(\lambda) = (\lambda - \lambda_1)m_U(\lambda). \tag{2}$$

In particular, A satisfies (i*) if and only if U does. But for $t = 1$, B is upper triangular if and only if U is. Hence, by induction, (i*) is equivalent to A being similar to an upper triangular matrix (details please!).

To show that (i*) *and* (ii*) *imply diagonalizability*, we choose t as big as possible and use the following similarity formula (to be checked by you):

$$\begin{bmatrix} I_t & -X \\ 0 & I_u \end{bmatrix} \begin{bmatrix} T & Y \\ 0 & U \end{bmatrix} \begin{bmatrix} I_t & X \\ 0 & I_u \end{bmatrix} = \begin{bmatrix} T & Z \\ 0 & U \end{bmatrix} \tag{3}$$

where X, Y, Z are $t \times u$ matrices. Y replaces the unspecific "$*$" of the last lesson, X remains to be picked, and Z equals $Y + TX - XU$.

In the present set-up, $T = \lambda_1 I_t$, $Z = Y + X(\lambda_1 I_u - U)$, and U may be assumed to be diagonal by induction — again using (2). If λ_1 occurred on the diagonal of U, there would be a W among $W_1, \ldots, W_u$ such that $(A - \lambda_1 I_n)W = V$ is a linear combination of $V_1, \ldots, V_t$ (see?). Since t is maximal, i.e. equal to the nullity of $(A - \lambda_1 I_n)$ or $(B - \lambda_1 I_n)$, we cannot have $V = 0$. Therefore the minimal polynomial of A acting on W (cf. preceding lesson) would be $m_{A|W}(\lambda) = (\lambda - \lambda_1)^2$. If (ii*) holds, this cannot happen, and hence $(\lambda_1 I_u - U)$ must be invertible. Thus, whatever Y might be, X can be chosen so as to make $Z = 0$, and we wind up with a diagonal matrix made up from T and U.

29. Jordan Blocks. For every integer $r > 0$, let J_r denote the $r \times r$-matrix obtained by augmenting the $(r-1)$-st identity matrix I_{r-1} by a trivial first column and last row (making $J_1 = 0$). Any matrix of the form $aI_s + J_s$, where a is a scalar, will be called the *Jordan block of degree s and value a*. Up to similarity, these matrices turn out to be the building blocks for *all* upper triangular matrices, as we now prove.

Lemma: Any upper triangular matrix A is similar to a "direct sum" of Jordan blocks $A_0, \ldots, A_l$, i.e. to a matrix of the form:

$$\begin{bmatrix} A_0 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & A_l \end{bmatrix} \tag{1}$$

Replacing A by $A + cI_n$, if necessary, we may assume that the first column of A is zero. Disregarding this column and the first row of A, we see an upper triangular $(n-1) \times (n-1)$ matrix B, which by induction we may assume to be a direct sum of Jordan blocks $B_0, \ldots, B_m$. At this point, the first column of our matrix is still 0, and the first row has the form $\big[0, R_0, \ldots, R_m\big]$, with the subrow R_k sitting above the block B_k. If an R_k is zero, the block B_k splits off as a direct summand of A (this is easiest to see for B_m), and our game is won by induction. Hence we may suppose that $R_k \neq 0$ for $k = 0, \ldots, m$. Note: whenever we permute the B_k by suitable similarities of A, the corresponding "obstructions" R_k are permuted accordingly.

After these preliminaries, the proof itself has two stages. Both of these use the similarity formula (3) of the preceding lesson, with $U = B_m$ and T , Y submatrices of A determined by that choice. The aim of the game is to choose X so as to make $Z = Y + TX - XU$ as trivial as possible. Note: if X and Y are zero below some r-th row, then so is Z.

Stage 1: We first reduce every obstruction to a single coordinate. To this avail, we restrict the $t \times u$ submatrices X, Y, Z to being zero below the first row. Since T has trivial first column, this makes $TX = 0$ and simplifies Z to $Y - XU$. Therefore the first row of $Y - Z = XB_m$ can be made equal to anything in the row-space of B_m.

If B_m is invertible, Z can be made zero, and the proof is finished. If not, then $B_m = J_u$, and the first row of Z can be reduced to $R_m = \big[a_m, 0, \ldots, 0\big]$, with $a_m \neq 0$. As every Jordan block occurring in B can be permuted to last place and hence play the role of B_m, we conclude that $B_k = J_{u(k)}$ for all $0 \leq k \leq m$. The corresponding obstructions turn the first row of A into $\big[0, R_0, \ldots, R_m\big]$, where each R_k has a single non-zero component a_k at the far left.

Stage 2: We shall annihilate R_m by using an alliance of B_0 and R_0. For this trick to work we need $u(0) \geq u(m)$, so let us ensure this by first interchanging the blocks B_0 and B_m if necessary. Then we multiply the first row of A by $1/a_0$ (and its trivial first column by a_0) to make $a_0 = 1$. This "fusion" of R_0 and $B_0 = J_{u(0)}$ produces a new Jordan block J_r in the upper left corner of the transformed A, with $r = u(0) + 1$ strictly greater than $s = u(m)$ — unless $m = 0$, in which case there is nothing more to do.

Now comes the crux: finding an X to kill R_m. Imagine a matrix H_{rs} made up of the first s columns of the transpose J_r^T, and observe what happens to it when we multiply either by J_r on the left or by J_s on the right: its subdiagonal 1's are pushed upward into the diagonal by the first, forward into the diagonal by the second, of these operations. However, the second operation leaves a zero in the upper left corner of the diagonal. Therefore, if $r > s$, we obtain the following equation of $r \times s$ matrices:

$$J_r H_{rs} - H_{rs} J_s = L_{rs}, \tag{2}$$

where L_{rs} is the $r \times s$ matrix which is zero except for a single 1 in the upper left. The $(n-s) \times s$ matrices X, Y, Z to be used in our similarity formula will be zero below row r. Denoting the first r rows of X by X' (and similarly for Y and Z), we clearly get $Z' = Y' + J_r X' - X' J_s$. Since Y is zero except for the lone scalar a_m, we have $Y' = a_m L_{rs}$. Using (2), this is easily annihilated by putting $X' = -a_m H_{rs}$.

30. Higher Determinants. Determinants of $n \times n$ matrices become less and less computable as n increases, and yet we cannot simply ignore them. Even if they cannot be evaluated exactly, they are sometimes the most effective tool for drawing certain approximate or qualitative inferences. Usually it is enough to know that they exist and have the properties summarized in the following theorem, which will be discussed (but not proved) below.

Theorem : For every n there is a function "det" which assigns to every $n \times n$-matrix A a scalar $\det A$ *in such a way that*

$$\det AB = \det A \cdot \det B, \qquad \text{and} \qquad \det C = d_1 \cdots d_n, \tag{1}$$

for C triangular with diagonal entries $d_1, \ldots, d_n$.

It is pretty clear that there will be at most one such function (per given n), because elementary matrices (being triangular) have their determinants specified by the theorem, while other matrices are products of elementary ones. But why should it exist at all for $n > 3$? Complex multiplication does not generalize to $n = 3$, the cross-product exists *only* for $n = 3$ — some things apply only to certain dimensions. Most beginners try to compute a 4×4 determinant by the "diagonal rule" (multiplying matrix entries along oblique lines, n slanted downwards and n slanted upwards). This happens to work fine for $n = 2, 3$ but no farther. To get results consistent with the theorem, a 4×4 matrix $A = (a_{ij})$ needs to have the determinant $D(A) =$

$$\begin{array}{l} +a_{11}a_{22}a_{33}a_{44} - a_{11}a_{22}a_{34}a_{43} - a_{11}a_{23}a_{32}a_{44} + a_{11}a_{23}a_{34}a_{42} + a_{11}a_{24}a_{32}a_{43} - a_{11}a_{24}a_{33}a_{42} \\ -a_{12}a_{21}a_{33}a_{44} + a_{12}a_{21}a_{34}a_{43} + a_{12}a_{23}a_{31}a_{44} - a_{12}a_{23}a_{34}a_{41} - a_{12}a_{24}a_{31}a_{43} + a_{12}a_{24}a_{33}a_{41} \\ +a_{13}a_{21}a_{32}a_{44} - a_{13}a_{21}a_{34}a_{42} - a_{13}a_{22}a_{31}a_{44} + a_{13}a_{22}a_{34}a_{41} + a_{13}a_{24}a_{31}a_{42} - a_{13}a_{24}a_{32}a_{41} \\ -a_{14}a_{21}a_{32}a_{43} + a_{14}a_{21}a_{33}a_{42} + a_{14}a_{22}a_{31}a_{43} - a_{14}a_{22}a_{33}a_{41} - a_{14}a_{23}a_{31}a_{42} + a_{14}a_{23}a_{32}a_{41}, \end{array} \tag{2}$$

assigned to it. We temporarily call this expression $D(A)$ instead of $\det A$ because it is still on probation. It is composed of 24 terms $\pm a_{1i_1}a_{2i_2}a_{3i_3}a_{4i_4}$, corresponding to the 24 permutations (i_1, i_2, i_3, i_4) of the indices $(1, 2, 3, 4)$. Each permutation can be reached from $(1, 2, 3, 4)$ by a sequence of "swaps" (exchanging 2 indices), and the sign $\pm$ indicates whether the particular (i_1, i_2, i_3, i_4) requires an even or an odd number of these. There are many ways, some shorter than others, of effecting a given permutation by a number of swaps; so it is unclear whether this is really a legitimate definition of our choice of signs.

If you study (2) attentively, you will find that each of its 4 lines can be rewritten as $a_{1j} \det A_{1j}$, giving

$$\Delta_1(A) = a_{11} \det A_{11} - a_{12} \det A_{12} + a_{13} \det A_{13} - a_{14} \det A_{14}, \tag{3}$$

where A_{1j} denotes the 3×3 submatrix obtained from A by crossing out row 1 and column j. If you look even harder, you will see that there are analogous "expansions" $\Delta_2(A)$, $\Delta_3(A)$, $\Delta_4(A)$ by rows 2, 3, 4. Of course we claim to have made up all $\Delta_i(A)$ so as to equal $D(A)$, but some back-benchers say that they do not quite see this — and their ranks will swell as n increases.

We want to argue as in lesson 12: if E is the matrix of an elementary row operation not involving row 1, we can use the expansion (3) to show that $\Delta_1(EA) = D(E)\Delta_1(A)$. Indeed, E induces simultaneous elementary row operations on the 3×3 matrices A_{1j} whose determinants appear in the expansion. Unfortunately we would also need $\Delta_2(A)$, etc., in order to cover all rows. If we could get everybody to agree that all $\Delta_i(A)$ equal $D(A)$, we would conclude that $D(EA) = D(E)D(A)$ for all elementary E, and arrive at the Theorem, with $\det A = D(A)$, by the usual reasoning.

So there's the rub: proving that $\Delta_i(A) = D(A)$, for general n. If you work down from $D(A)$, you must first firm up its definition by justifying the sign-rule via a detour through permutations and swaps. If you start with $\Delta_i(A)$, you can build on your knowledge of $(n-1) \times (n-1)$ determinants to verify certain formal properties of $\Delta_i(A)$ considered as a function of the n columns (it is an "alternating n-form"). Then it is fairly routine to check that $\Delta_i(AB) = D(A)\Delta_i(B)$. Here $D(A)$ needs no prior justification because it is the *result*, not the basis, of a computation. The rest is a breeze.

F. The Hessenberg Process. Through much of this course we have been interested in replacing a given square matrix A by a *similar* one which was (ideally) diagonal or had some other standard form revealing its true nature. Even if we settled for just a triangular form, we would need to start with at least one eigenvector of A, which is usually hard to find.

On this page we shall describe a process which simply makes A easier to handle, without attempting to analyse it. Like the elimination method it works by easy instalments, one column at a time. However, since it involves similarity instead of just left multiplication, it is inevitably more complex. Its end-product is a matrix similar to A but with a lot of zeroes in the lower left corner.

Let us say that A is *subtriangular* (or a *"Hessenberg matrix"*) if its i,j-entry is zero whenever $i > j+1$. We shall see how to transform *any $n \times n$ matrix into a similar subtriangular matrix* by a succession of $n-2$ conjugations. The s-th of these has the form:

$$\begin{bmatrix} I_s & 0 \\ 0 & H \end{bmatrix} \begin{bmatrix} A' & Y \\ X & A'' \end{bmatrix} \begin{bmatrix} I_s & 0 \\ 0 & H^{-1} \end{bmatrix} = \begin{bmatrix} A' & YH^{-1} \\ HX & HA''H^{-1} \end{bmatrix} \tag{1}$$

where the matrices A', X, and H have dimensions $s \times s$, $(n-s) \times s$, and $(n-s) \times (n-s)$, respectively (we do not care about the others). At the beginning of the s-th step, all the columns $C_1, \ldots, C_s$ of the submatrix X are zero, except (possibly) the last one C_s. Then, if $n - s \geq 2$, we can choose an H such that $HC_s = (c, 0, \ldots, 0)^T$ (see?), so that the only non-zero entry (if any) of HX is in the upper right corner. As we do this to X, formula (1) shows that the upper left A' stays put. This trick will be used to dig a cave of zeroes into the lower left side of the matrix, as follows.

Start with $s = 1$ and X being the portion of the first column below the first row. Using (1) turn the first column into 0 except for the first two entries. Then go on to $s = 2$, with X being the part of the first two columns below the second row. At every stage, the new X has one more column but one less row than its predecessor: as the cave extends farther into the matrix, its ceiling gets lower. The process stops after $n-1$ steps.

For illustration, we show the situation for $n = 5$ after two steps. Here $s = 3$, and you are preparing to change the (x, y) into a $(\star, 0)$:

$$\begin{bmatrix} \star & \star & \star & * & * \\ \star & \star & \star & * & * \\ 0 & \star & \star & * & * \\ 0 & 0 & x & \bullet & \bullet \\ 0 & 0 & y & \bullet & \bullet \end{bmatrix}. \tag{2}$$

At every new value of s, you have already performed $s-1$ such operations and created a subtriangular $s \times s$-matrix A' with an $(n-s) \times s$-block X below it, which is zero except for its last column.

The submatrix H used in every one of these steps only needs to be invertible and satisfy $HC_s = (c, 0, \ldots, 0)^T$. Hence you could use either a sequence of row operations (as in Gaussian elimination) or a suitable reflection (as in the QR-factorization). In either case you must remember to apply the corresponding *inverse* operation to the columns, in order to get the third factor in (1). Using reflections would have the advantage of yielding an orthogonal similarity. Thus, we see that *every square matrix is* orthogonally *similar to a subtriangular matrix.* Not only that: we have an easy systematic process for achieving this similarity.

Now, if A is symmetric, so is $B = Q^T AQ$ for every orthogonal Q. Hence, if b_{ij} is zero for $i > j+1$, it will also be zero whenever $j > i+1$. Therefore its only non-zero entries (if any) are *on* the diagonal or just on either side of it. Such a matrix is called *tridiagonal.* So we have shown how *every symmetric matrix can be transformed into a tridiagonal matrix by orthogonal similarities.* Of course, this might not be good enough. But even if you must obtain the actual *diagonal* form (cf. lesson 26), the algorithm described above is a useful preparation for the final assault: there will be fewer off-diagonal entries to fight.