An Initiation into Linear Algebra

Any fool can know. The point is to understand.
Albert Einstein

Dr. Ramakalyan Ayyagari ©

Professor, Dept of ICE

NIT Tiruchirappalli, 620 015, INDIA

rkalyn@nitt.edu

Contents

1 An Initiation into Linear Algebra

Linear Algebra is an old course. It was in the domain of mathematicians for a long time, with sporadic applications in, and hence some contributions from, engineers especially in the computer sciences and electrical engineering departments. However, the ideas of linear algebra presently have a huge impact on a host of engineering topics and there is a rush outside. Consequently, any interested learner has multiple options of video lectures, certificate courses, much better written artilcles and textbooks. While the so-called modern control is built heavily on the concepts of linear algebra, and is being taught in engineering schools over the last three decades, or even more, an *appropriate* introduction and guidance has been away and there is every need for the student to learn it ab initio. This short primer (of 40 odd pages) has two parts. The first part, a slightly longer one, develops the ideas from scratch; the theme is that all of us know the subject informally, but it is time to erase possible pitfalls and learn (which is different from mere knowing) it more formally. The second part, perhaps very short in contrast, attempts to align the material from the first part with formal definitions and algorithms, through carefully chosen arguments and examples. We avoid theorems and proofs for want of space; nevertheless they are readily available elsewhere for an initiated student.

Routinely, we learn about eigenvalues (the key concept of linear algebra) through bland definitions: non-trivial solutions, $\mathbf{t_i} \neq 0$ of the linear system$(A - \lambda_i I)\mathbf{t_i} = 0$ are called eigenvectors, and the scalars λ_i are called eigenvalues. We do not gain much insight into the concept by reading this. It appears that the intuition takes a back seat all of a sudden, without prior notice and observations, and the manipulative skills to solve numericals take the driver's seat. In this note we try to present the concepts via insightful questions first and then develop more formal arguments which eventually align with the standard theories. We assume that the reader is familiar with alphabet of matrices – elements, rows, columns, vectors, addition and (non-commutative) multiplication of matrices. A formal course on linear algebra begins with vector spaces where the eigenvalues and eigenvectors of matrices play a key role. We will begin with building a bridge to linear algebra.

1.1 Matrices in Day-to-Day Matters

Let us begin with a familiar problem disguised as a game as follows. An honest, as well as intelligent, storekeeper at an Institute of Mathematics sells exotic souvenir items $1, 2, \cdots, n$, but he insists that every buyer should buy at least two different items; to the buyer he does not disclose the individual prices x_i, but bills her exact amount b_i; however, the buyer is not permitted to visit the store for a second time. A curious student buyer, together with her friends, brainstorms about determining the prices and develops the following strategy, which the reader too would have thought.

First, let us assume there is only one item in the shop. This leads to

$$a_{11}x_1 = b_1$$

where a_{11} is the number of units purchased and b_1 is the amount paid, from which the price x_1 of the item can be readily determined, provided $a_{11} \neq 0$.

Next, if there are two items, she would first model her transaction as follows:

$$a_{11}x_1 + a_{12}x_2 = b_1 \tag{1}$$

where a_{11} and a_{12} are the number of units of item #1 and item #2 respectively, and b_1 is the cumulative amount paid to the storekeeper. Then, she would encourage one of her friends to make a different transaction. Quite obviously, the friend cannot repeat the first transaction – this leads them nowhere; the second transaction:

$$a_{21}x_1 + a_{22}x_2 = b_2 \tag{2}$$

where a_{2i} is the number of units of item #i purchased by the friend. This model of transaction must be *independent* of the first transaction in the sense that the well-known condtion:

$$\frac{a_{11}}{a_{21}} \neq \frac{a_{12}}{a_{22}} \quad \text{or, equivalently, } a_{11}a_{22} - a_{12}a_{21} \neq 0 \tag{3}$$

holds, i.e., the friend not only cannot repeat the original transaction, but cannot make any *proportionate* transaction. This condition generalizes, in a strong sense, the trivial requirement of $a_{11} \neq 0$ in the simplest scalar case mentioned above. Indeed, a -ve sign appears from nowhere, but this helps us choose coefficients appropriately. We will make a more detailed commentary on this most important concept of independence on the fly.

A natural question is – is there a clever way of choosing a_{21} and a_{22}? Without much effort one can simply think of $a_{21} = a_{11}$ (not violating eqn. (3)) so that the difference $b_1 - b_2$ in the bills is proportional to x_2; with x_2 known, it is easy to compute x_1. In terms of matrices and vectors, the above two transactions can be put together as follows:

$$\begin{bmatrix} a_{11} & a_{12} \\ 0 & a'_{22} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} b_1 \\ b'_2 \end{bmatrix} \tag{4}$$

where $a'_{22} = a_{22} - a_{21}$ and $b'_2 = b_2 - b_1$. More generally, if $a_{21} = k \cdot a_{11}$, where k is a scalar, then $a'_{22} = a_{22} - k \cdot a_{12}$ and $b'_2 = b_2 - k \cdot b_1$; this gives more freedom to the second buyer.

If the elements a_{ij} of the matrix are such that

$$\frac{a_{11}}{a_{21}} = \frac{a_{12}}{a_{22}} = k \left(= \frac{b_1}{b_2} \right) \tag{5}$$

then $r_1 - k \cdot r_2 = 0$, where r_i indicates the i^{th} row of the matrix, and this is an indication of proportionality or *dependence* of the rows. We look for independence.

Let us generalize this a little more now. For a unique solution, first, we wish that the 2^{nd} row is *not* proportional to the first row, as in the 2×2 matrix case. Extending the argument logically, in a $n \times n$ case, what we expect is the following:

$$\begin{aligned} \text{either row 1} &\not\propto \text{row 2} \\ \vee \text{ row 1} &\not\propto \text{row 3} \\ &\vdots \\ \vee \text{ row 1} &\not\propto \text{row n} \end{aligned} \tag{6}$$

where $\vee$ stands for the logical disjunction, "or." If there is at least one proportionality satisfied, then the uniqueness of the problem would be shattered as we establish a few steps ahead. Put together, we have that

$$r_1 \neq \bigvee_{i=2}^{n} \alpha_i r_i \tag{7}$$

or, translated mathematically,

$$\exists \alpha_i \, : \, \alpha_1 r_1 - \sum_{i=2}^{n} \alpha_i r_i \neq 0 \tag{8}$$

where α_is are proportionality constants such that $\alpha_i \neq 0$ indicates proportionality and, the complement $\alpha_i = 0$ indicates no proportionality. Accordingly, the entire expression is treated as a sum of products in the Boolean sense.

Formally, we look at an algebraic addition of the scaled rows of the matrix:

$$\sum_{i=1}^{2 \leq k \leq n} \alpha_i r_i = \alpha_1 \cdot r_1 + \alpha_2 \cdot r_2 + \cdots + \alpha_k \cdot r_k \tag{9}$$

known as a linear combination of the rows. The rows r_1 to r_k are said be linearly independent, generally written as *l.i.*, if the linear combination is such that $\sum \alpha_i r_i = 0$ if and only if $\alpha_1 = \alpha_2 = \cdots \alpha_k = 0$; otherwise, the k rows are linearly dependent. More on this a little later.

The idea is now clear – for the number of items n, it is necessary and sufficient[1] to have n friends making n transactions which can be packed as follows:

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} & \cdots & a_{1n} \\ 0 & a'_{22} & a'_{23} & \cdots & a'_{2n} \\ 0 & 0 & a''_{33} & \cdots & a''_{3n} \\ \vdots & & & \ddots & \\ 0 & 0 & 0 & \cdots & a_{nn}^{(n-1)} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} b_1 \\ b'_2 \\ b''_3 \\ \vdots \\ b_n^{(n-1)} \end{bmatrix} \tag{10}$$

It is necessary because if the transactions is one less, then the price of one of the items remains indeterminate, and sufficient because one more transaction is redundant. Needless to say that none of the *pivot* elements, $a_{11}, a'_{22}, \cdots a_{nn}^{(n-1)}$, along the principal diagonal is allowed to be a zero. There is a slight difference between diagonal elements and pivots – a diagonal element a_{ii} becomes a pivot if the elements of i^{th} row, $a_{i1} \cdots a_{i,i-1}$, are reduced to zero by way of performing "scaling-and-adding" operations, a.k.a. linear combinations, on the rows.

In the general case, where the transactions can be arbitrary with more freedom to the buyers, what is essential is to establish the *independence*, akin to eqn. (3), among the n transactions, which in turn would ensure non-zero pivots in every row of the matrix.

First things first, as n gets bigger the number of possible transactions explodes combinatorially; we should be able to arrive at a method to decide whether a given set of n arbitrary transactions are independent. We wish to ask:

? Is the matrix "qualified" to offer a unique solution

Because the pivotal positions are unique, it follows that the number of pivots, which is the same as the number of independent rows in A, is also uniquely determined by the entries in A. This integer is called the *rank of A*. Thus, the answer to the above question is yes if the matrix has *full rank*, n. We continue our investigation by asking: How do we detect if the rank $< n$? We will do this first with $n = 3$ and then generalize it to n.

We will quickly study the following matrices which are indicative and illustrative of possible transactions. For convenience, let us assume that the

[1] literally

store has 3 units of each of the 3 items.

$$A_1 = \begin{bmatrix} 3 & 1 & 1 \\ 0 & 2 & 1 \\ 0 & 0 & 1 \end{bmatrix}, \quad A_2 = \begin{bmatrix} 3 & 2 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix}, \quad A_3 = \begin{bmatrix} 1 & 2 & 1 \\ 1 & 1 & 2 \\ 1 & 0 & 0 \end{bmatrix},$$

$$A_4 = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}, \quad A_5 = \begin{bmatrix} 3 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 1 & 1 \end{bmatrix}$$

We get an insight with the following observations:

1. Matrices A_1 and A_2 present straight cases.

2. A_3 is not so straight, yet it gives us important clues. First, x_1 can be computed straight. Algebraically, yes making it formal gradually, the information in A_3 may be recast as

$$A_3' = \begin{bmatrix} 0 & 2 & 1 \\ 0 & 1 & 2 \\ 1 & 0 & 0 \end{bmatrix}$$

by replacing the first row with the difference between first and third rows, and similarly replacing the second row with the linear combination $r_2 - r_3$. Having solved for x_1 uniquely using the third row, we should be able to focus on the top right square block – rows 1 and 2, and columns 2 and 3, obtained by deleting the first column and the last row. We recognize that eqn. (3) stands satisfied, and thus we get x_2 and x_3 uniquely.

3. With the experience gained so far, proportionality is way too obvious (though we are yet to formally prove it) and matrix A_4 does not serve any purpose and we discard it.

4. Matrix A_5 may tend to fool us with its deceptive similarity to A_3. A step deeper, we readily see that two of the buyers (rows 2 and 3) essentially make the same purchase, and hence our problem remains unsolved as a result of the vanishing pivots.

Let us now take any 3×3 matrix:

$$A = [a_{ij}], \quad i, j = 1, 2, 3$$

Suppose buyer #1 completed his transaction: $\Sigma a_{1j} x_j = b_1$. Then, the strategy for the other two buyers is such that any two of x_i are uniquely solvable so that substituting them in the first transaction yields the third x_i. This can be done in three different ways, for a 3×3 case –

1. Irrespective of a_{21} and a_{31}, i.e., quantities of item 1 bought by buyers #2 and #3, ensure that $a_{22}a_{33} - a_{23}a_{32} \neq 0$ so that x_2 and x_3 may be determined from the 2^{nd} *and* (logical conjunction $\wedge$) 3^{rd} transactions followed by x_1 from the first transaction with a_{11} as the pivot, i.e.,

$$(a_{11} \neq 0) \wedge (a_{22}a_{33} - a_{23}a_{32} \neq 0) \tag{11}$$

or ($\vee$)

2. Irrespective of a_{22} and a_{32}, ensure that $a_{21}a_{33} - a_{23}a_{31} \neq 0$ so that x_1 *and* x_3 may be determined, while a_{12} is the pivot to determine x_2, i.e.,

$$(a_{12} \neq 0) \wedge (a_{21}a_{33} - a_{23}a_{31} \neq 0) \tag{12}$$

or ($\vee$)

3. Irrespective of a_{23} and a_{33}, ensure that $a_{21}a_{32} - a_{22}a_{31} \neq 0$ so that x_1 *and* x_2 may be determined, while a_{13} serves as the pivot to compute x_3, i.e.,

$$(a_{13} \neq 0) \wedge (a_{21}a_{32} - a_{22}a_{31} \neq 0) \tag{13}$$

We need to look into stiching these three statements logically together, i.e., identify the disjunction

$$\begin{aligned} \bigvee \{ &(a_{11} \neq 0) \wedge (a_{22}a_{33} - a_{23}a_{32} \neq 0)\,, \\ &(a_{12} \neq 0) \wedge (a_{21}a_{33} - a_{23}a_{31} \neq 0)\,, \\ &(a_{13} \neq 0) \wedge (a_{21}a_{32} - a_{22}a_{31} \neq 0)\} \end{aligned} \tag{14}$$

so that we are assured of a unique solution. We do this by looking at the contradiction that does not lead to a solution. Suppose that the first row is a linear combination of the second and the third rows, i.e., let there exist scalars $\alpha_1 \neq 0$ and $\alpha_2 \neq 0$ such that

$$a_{11} = \alpha_1 a_{21} + \alpha_2 a_{31}, \quad a_{12} = \alpha_1 a_{22} + \alpha_2 a_{32}, \quad \text{and} \quad a_{13} = \alpha_1 a_{23} + \alpha_2 a_{33}$$

Substituting this in the three logical statements above and separating out the terms multiplied by α_1 and α_2, we get

$$\begin{aligned} S1: &\quad \alpha_1 (a_{21}a_{22}a_{33} - a_{21}a_{23}a_{32}) + \alpha_2 (a_{31}a_{22}a_{33} - a_{31}a_{23}a_{32}) \\ S2: &\quad \alpha_1 (a_{22}a_{21}a_{33} - a_{22}a_{23}a_{31}) + \alpha_2 (a_{32}a_{21}a_{33} - a_{32}a_{23}a_{31}) \\ S3: &\quad \alpha_1 (a_{23}a_{21}a_{32} - a_{23}a_{22}a_{31}) + \alpha_2 (a_{33}a_{21}a_{32} - a_{33}a_{22}a_{31}) \end{aligned} \tag{15}$$

We seamlessly observe that

$$S1 - S2 + S3 \;=\; 0 \;\; \forall\, \alpha_1, \alpha_2 \neq 0 \tag{16}$$

and the transactions are dependent. In other words, the sum of products

$$\begin{aligned}\Delta_3 &= a_{11} \cdot (a_{22}a_{33} - a_{23}a_{32}) \\ &\quad - a_{12} \cdot (a_{21}a_{33} - a_{23}a_{31}) \\ &\quad\quad + a_{13} \cdot (a_{21}a_{32} - a_{22}a_{31}) \qquad (17)\end{aligned}$$

is an indication of the proportionality among three transactions – the three rows of the matrix are independent $\Leftrightarrow$ $\Delta_3 \neq 0$. This is a generalization to eqn. (3), and this can be extended to accommodate any number of transactions.

The above argument may be equivalently presented from the point of view of items as well.

1. Suppose buyer #1 purchased item #1; then the other two buyers should make their purchases in such a way that $a_{22}a_{33} - a_{23}a_{32} \neq 0$; this is same as the first piece of the previous argument.

2. Suppose, now, buyer #2 purchases item #1, then we find that $a_{12}a_{33} - a_{13}a_{32} \neq 0$, and

3. if buyer #3 purchases item #1, then we find that $a_{12}a_{23} - a_{13}a_{22} \neq 0$.

And, we may now work out a bit routinely to figure that

$$\begin{aligned}\Delta_3 &= a_{11} \cdot (a_{22}a_{33} - a_{23}a_{32}) \\ &\quad - a_{21} \cdot (a_{12}a_{33} - a_{13}a_{32}) \\ &\quad\quad + a_{31} \cdot (a_{12}a_{23} - a_{13}a_{22}) \qquad (18)\end{aligned}$$

also.

Thus, we may check for the row independence, as well as column independence. An avid reader would notice two quick things – eqn. (17) and eqn. (18) both have 6 terms – 3 +ve and 3 -ve, and moreover, it is actually the same set of 6 products (containing all the 9 elements of the matrix) but seen starting with a different header. Morphing algebraically, Δ_3 can be computed in 6 equivalent ways – along the three rows, and along the three columns. For instance, along the second row,

$$\begin{aligned}\Delta_3 &= -[a_{21} \cdot (a_{12}a_{33} - a_{13}a_{32}) \\ &\quad - a_{22} \cdot (a_{11}a_{33} - a_{13}a_{31}) \\ &\quad\quad + a_{23} \cdot (a_{11}a_{32} - a_{12}a_{31})] \qquad (19)\end{aligned}$$

It is with a purpose, the sign convention, that we jumbled the triplets and brought out the -ve sign outside. Looking back at eqn. (3), since we are primarily interested in checking whether the expression evaluates to a zero or not from the ratios, we might equally likely consider the difference $a_{12}a_{21} - a_{11}a_{22}$. However, the earlier one has been universally accepted and the reader must have been familiar with similar cross-multiplication arithmetic in high school. In fact, from the second row point-of-view, the matrix may be rewritten as

$$\begin{bmatrix} \leftarrow & \text{row } 2 & \rightarrow \\ \leftarrow & \text{row } 3 & \rightarrow \\ \leftarrow & \text{row } 1 & \rightarrow \end{bmatrix} \tag{20}$$

to justify the sign convention – products downwards are +ve and upwards are -ve.

Henceforth, we call the expression in eqn. (3) as Δ_2. Since a non-zero $\Delta_1\,(= a_{11})$, or Δ_2, or Δ_3 is necessary for a unique solution, we call them **determinants** of the matrices, denoted as $|A|$. We need to stick to such details because in addition to checking whether it is zero or not, its magnitude as well as sign play an important role in subsequent theory and applications. One may, out of sheer curiosity, refer to a dictionary to find that

determinant (Oxford Dictionary): a factor which *decisively* affects the nature or outcome of something – e.g., pure force of will was the main determinant of his success.

Next, we take up the general case of a $n \times n$ matrix and develop a uniform formula that is applicable for any square matrix. For every element a_{ij} of the matrix (of size $\geq 2 \times 2$), we define

$$\text{Minor } M_{ij} \triangleq \text{determinant of the lower order matrix obtained by deleting } i^{th} \text{ row and } j^{th} \text{ column}$$

Going by this,

$$\Delta_2 = a_{11}M_{11} - a_{12}M_{12}$$

$$\Delta_3 = a_{11}M_{11} - a_{12}M_{12} + a_{13}M_{13}$$

also.

Notice that we have been walking along the first row. We might also observe that, in the case of 2×2 matrix, if we walk along second row we get the negative of $\mathcal{D}_2$. In case of $\mathcal{D}_3$ we have already shown that walking along 2^{nd} row needs changes in the sign, and the original sign pattern restores for the third row, i.e.,

$$\begin{aligned} \Delta_3 &= -a_{21}M_{21} + a_{22}M_{22} - a_{23}M_{23} \\ &= +a_{31}M_{31} - a_{32}M_{32} + a_{33}M_{33} \end{aligned}$$

Accordingly we make the following additional definition:

$$\text{Co-factor, } C_{ij} \triangleq (-1)^{i+j} M_{ij} \tag{21}$$

so that walked along any i^{th} row for any $n \times n$ matrix,

$$\Delta_n = \sum_{j=1}^{n} a_{ji} C_{ji} \tag{22}$$

This is the well-known Laplacian expansion of a determinant. Here we get two interesting observations:

1. We can walk along columns also. For instance,

$$\Delta_2 = a_{11}C_{11} + a_{21}C_{21} = a_{12}C_{12} + a_{22}C_{22}$$

   is readily found to conform. when $n = 3$, it is indeed meaningful to have the strategic statements such as: irrespective of the purchase of buyer #i, the purchase by the remaining two buyers should be able to yield the prices of two items. Accordingly, in general, walking along any j^{th} column,

$$\Delta_n = \sum_{j=1}^{n} a_{ij} C_{ij} \tag{23}$$

2. Since the basic idea is that all the n buyers make n *independent* transactions, walking across rows/columns yields no result, equivalent to dependence. In other words,

$$a_{11}C_{12} + a_{12}C_{22} + a_{13}C_{32} = 0 \tag{24}$$

   and so are other similar sum of products $\sum a_{ij}C_{pq}$, $i \neq p$ and $j \neq q$.

These two observations may be neatly expressed in the following matrix equation:

$$[a_{ij}]\,[C_{ij}]^T = \Delta_n \cdot I_n = [C_{ij}]\,[a_{ij}]^T \tag{25}$$

where the transposition of one of the matrices is obvious, and I_n is the identity matrix of size n.

We define adjoint of A,

$$adj(A) \triangleq [C_{ij}]^T \tag{26}$$

leading to

$$A \cdot \frac{adj(A)}{\Delta_n} = I_n$$

from which we readily understand that the *inverse* of a matrix is

$$A^{-1} \triangleq \frac{adj(A)}{\Delta_n} \tag{27}$$

Evidently, the inverse A^{-1} exists *if and only if* all the rows, and now all the columns as well, are independent; in other words,

$$displaystyle \exists A^{-1} \Leftrightarrow \Delta(A) \neq 0$$

We will close this section with an example.

Example 1: Let

$$A = \begin{bmatrix} 3 & 0 \\ 4 & 1 \end{bmatrix}$$

The matrix of minors is

$$M = \begin{bmatrix} 1 & 4 \\ 0 & 3 \end{bmatrix}$$

The matrix of cofactors is

$$C = \begin{bmatrix} 1 & -4 \\ 0 & 3 \end{bmatrix}$$

Hence, the adjoint is

$$adj(A) = \begin{bmatrix} 1 & 0 \\ -4 & 3 \end{bmatrix}$$

from which determinant is

$$\Delta_2 = 3 \cdot 1 + 0 \cdot (-4) \; = \; 3 \; = \; 4 \cdot 0 + 1 \cdot 3$$

expanded along the two rows; column-wise expansion would also yield identical result, and the reader may quickly verify. and, the inverse of A is

$$A^{-1} = \frac{1}{3} \begin{bmatrix} 1 & 0 \\ -4 & 3 \end{bmatrix}$$

One may readily verify that $A \cdot A^{-1} = A^{-1} \cdot A = I_2$

To sum up, a set of n rows is *l.i.* if and only if the determinant $\Delta_n \neq 0$. In the most ideal case independent rows show up readily in a diagonal matrix, with non-zero diagnonal elements. In a more practical case we see a triangular matrix. Arbitrary matrices need to be investigated further by computing the determinant. Interestingly, there are innumerable possibilities in the placement of the elements of an arbitrary $n \times n$ matrix, and each possibility has a thick story; for instance, the reader must have been already familiar with certain symmetric matrices whose inverses are the matrices themselves, i.e., $\exists A : A^T = A$ and $A^{-1} = A$ or $A^2 = I$, and such matrices are known as **orthonormal** if the elements are real, or **unitary** matrices, more generally, if the elements are complex numbers.

At this juncture, we make some subtle and not too intuitive observations that help us align with the mainstream linear algebra.

1. First, the store has n items; this can be one (of course, we have no fun), or two, or a large finite number N, or even infinity. In a way, if there is something very special about these n items (say, exotic ice-cream flavours[2] not available elsewhere) from that store then we *attribute* the store with the number n of items.

2. Secondly, in some sense one can think of these n items as distinct (from each other) but could be 'combined' to make more items; one can have one scoop of vanilla flavour and two scoops of chocolate flavour and enjoy the combination of ice-creams. This permits us to have an infinitely big store where every item is either one of the n base items, or combinations of the base items; we consider that combinations of combinations are still combinations of the base items. The combinations need a lot of creativity, obviously.

3. Next, if we look at the combined items little more closely, we see that they are either homogeneous, e.g., 'two' scoops of chocolate, or two or more flavours *added* to each other, or both; for instance, one can take away a combination like $x_1 + 2x_2 + \sqrt{5}x_4$. Thus, we do speak about what are known as *linear* combinations.

4. Next, there is certainly a possibility of the store getting empty by removing all the items, and furthermore there could be orders pending arrival of fresh stocks.

5. Next, we can look at a "transformation," say, from the quantities or rates of items purchased to the bills paid, and vice-versa; in fact, we have begun our story with this, visualizing a simple linear system of equations where the individual rates x_i have been transformed by the store keeper to the bills paid by each of the friends, and the friends figuring out a transformation the other way round.

[2]Baskin-Robbins is known for its "31-flavour" slogan, intending to serve customers with a different flavour every day of any month.

6. Last, but not least, we are looking at general structures. Needless to say, once we are able to establish a clean transformation, e.g., n friends making n independent transactions, we have our problem formulated. When we look around we see a multitude of problems formulated in very much the same way, and if we have an algorithm to solve this *structure* called linear system of equations, then we have done our job. More specifically, we are equipped with tools such as vectors and matrices, and a set of well-defined operations like matrix-matrix products, determinants, inverses and so on.

With these observations, we may now look at a formal definition of a linear vector space as follows. This helps us immensely in quantifying the aforementioned features – clean transformation among the data gathered, and the algorithm to infer the unknown for a general, usually larger, n. We purporsefully avoid emphasizing such things as symbols or names or applications; nevertheless, properties derived exclusively from the structure will hold for anything that has the same structure. Most appropriately, we say all such problems have a *vector-space* structure.

1.1.1 Formal Definition of a Vector Space

We will first briefly introduce one of the three major structures of algebra which underlie our operations, theorems, and results in vector spaces.

Field $\mathcal{F}$

A field is any number system in which, roughly speaking, we can add, subtract, multiply, and divide according to the usual laws of arithmetic.

Typical examples include the reals $\Re$, the complex numbers $\mathcal{C}$, and a finite set of prime numbers. One may readily realize that the set integers cannot make a field.

The other two structures are called Groups and Rings and are covered in detail in courses in abstract algebra. Interested reader may refer to the references.

Vector Space over $\mathcal{F}$

A vector space $\mathcal{V}^n(+, \cdot)$ over a field $\mathcal{F}$ is a non-empty set

$$\{v_1, v_2, \cdots, v_n, v_{n+1}, \cdots\}$$

of mathematical objects called **vectors**, together with

- a rule '+' for adding a pair of vectors, say v_i and v_j to produce $v_k = v_i + v_j$ which also belongs to the same space $\mathcal{V}^n$, and
- a rule '$\cdot$' for scaling any vector v_i with a scalar $\alpha_i \in \mathcal{F}$ such that $\alpha_i \cdot v_i \in \mathcal{V}^n$.

Moreover, there must exist a vector $\bar{0}$ (read as the zero vector, or more appropriately the origin) in $\mathcal{V}^n$ as well as the vector $-v_i$, the negative of any vector $v_i \in \mathcal{V}^n$, satisfying the following axioms[3] [4]:

$\forall \alpha_0, \alpha_1, \alpha_2 \in \mathcal{F}$ and $\forall v_0, v_1, v_2, v_3 \in \mathcal{V}^n$

S1. $\alpha_0 \cdot (v_1 + v_2) = \alpha_0 \cdot v_1 + \alpha_0 \cdot v_2 \in \mathcal{V}^n$

S2. $(\alpha_1 + \alpha_2) \cdot v_0 = \alpha_1 \cdot v + \alpha_2 \cdot v_0 \in \mathcal{V}^n$

S3. $(\alpha_1 \cdot \alpha_2) \cdot v_0 = \alpha_1 \cdot (\alpha_2 \cdot v_0) \in \mathcal{V}^n$

S4. $\exists 1 \in \mathcal{F} \,:\, 1 \cdot v = v \in \mathcal{V}^n$

A1. $v_1 + v_2 = v_2 + v_1 \in \mathcal{V}^n$

A2. $(v_1 + v_2) + v_3 = v_1 + (v_2 + v_3) \in \mathcal{V}^n$

A3. $\exists \bar{0} \in \mathcal{V}^n \,:\, \bar{0} + v = v \in \mathcal{V}^n$

A4. $v + (-v) = \bar{0} \in \mathcal{V}^n$

Remark: Observe that '$\cdot$' operation just 'scales a vector'; that is why we enumerated the axioms as S1 to S4 ('S' for scaling) whereas the reader would have found M1 to M4 elsewhere. Multiplication of two vectors is not a part of the structure, though certain products, e.g., dot product, can be defined and they are treated as additional structure on the vector space.

Example 2: The set of integers $\mathcal{Z}$ over the field $\mathcal{F}$ of integers is a vector space. However, if the field is all real numbers the set $\mathcal{Z}$ is not a vector space. Typically, $\Re$ over $\Re$ qualifies as a vector space, so much so that we perform the two operations, without even acknowledging it as a *space*; we refer to it simply as the *real line*, with a 0 and a unit 1.

One can also look at the x-y axes: $\Re \times \Re$, familiar from childhood, as a vector space $\Re^2$ over $\Re$. In almost every practical application we get $\Re^n$ over $\Re$.

[3] The German mathematician Hermann Grassmann (1809–1877) is generally credited with first introducing the idea of a vector space in 1844. Unfortunately, his work was very difficult to read and did not receive the attention it deserved. Later, the Italian mathematician Giuseppe Peano (1858–1932), in his 1888 book *Calcolo Geometrico*, clarified Grassmann's work and laid down the axioms for a vector space as we know them today. Peano's axiomatic definition of a vector space also had very little influence for many years. It was only in 1918, after Hermann Weyl (1885–1955) repeated it in his book *Space, Time,Matter, an introduction to Einstein's general theory of relativity* the community accepted the axioms.

[4] An *axiom* is a statement that is taken to be true, to serve as a premise or starting point for further reasoning and arguments. The word comes from ancient Greek *axiōma*, meaning 'that which is thought-worthy or fit,' or 'that which commends itself as evident.'

With the two rules of scalar multiplication and (ordinary, well-known) addition we now formally define

$$\sum_{i=1}^{n} \alpha_i v_i = \alpha_1 v_1 + \alpha_2 v_2 + \cdots + \alpha_n v_n \tag{28}$$

as a **linear combination** of the vectors $v_i \in \mathcal{V}$ using the scalars $\alpha_i \in \mathcal{F}$, nicely put as *saxpy* for *scalar a x plus y* by Golub and Van Loan [6]. Following the axioms we understand that the result of such a linear combination is another vector in the same space. We may be tempted to examine this newly constructed vector in more details.

1. Suppose the linear combination results in a scaled version of any of the n vectors, i.e.,

   $$\alpha_1 v_1 + \alpha_2 v_2 + \cdots + \alpha_n v_n = \beta_i v_i, \quad \beta_i \in \mathcal{F}$$

   It is easy to see that

   $$v_i = \sum_{j \neq i} \alpha'_j v_j, \quad \text{not all } \alpha'_j = 0$$

   i.e., v_i is now a linear combination of the rest of the vectors. If such is the case, we say that the set of vectors $v_i \cdots, v_n$ is **linearly dependent**.

   It is interesting to note that, among the given set of n vectors, if there is the zero vector then the set is linearly dependent. The converse is not necessarily true.

2. On the other hand, in fact negating logically the above observation, we formally define a set of vectors $v_1, \cdots, v_n$ is **linearly independent** (*l.i.* in short)

   $$\alpha_1 v_1 + \cdots + \alpha_n v_n = 0 \;\Leftrightarrow\; \alpha_i \,(\in \mathcal{F}) = 0,\, i = 1, \cdots, n$$

   In other words, there is no linear relation among the vectors in the set, except the trivial one in which all the coefficients α_i are zero.

The following hypothetical example illustrates several ideas.

Example 3:

If we consider the 26-letter English alphabet, where in place of z the letter h repeats, we will not be able to access words containing z, and the additional h does not help us; our vocabulary is restricted.

Let us, for a moment, imagine of these 26 letters as vectors $\in \mathcal{A}^{26}$:

$$v_a = \begin{bmatrix} a \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, v_b = \begin{bmatrix} 0 \\ b \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \cdots, v_z = \begin{bmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ z \end{bmatrix}$$

It is easy to infer that this is a *l.i.* set – none of the vectors can be obtained as a linear combination of the others. If v_h replaces v_z we will have

$$0 \cdot v_a + 0 \cdot v_b + \cdots + \ 1 \cdot v_h \ + 0 \cdot v_i + \cdots \ + (-1) \cdot v_h = 0$$

i.e., there exist two non-zero scalars, $\alpha_h = 1$ and $\alpha_z = -1$, to prove that it a linearly dependent set.

Implicit in this example is a core idea. If the space has exactly n independent vectors and, hence, any other vector is a **unique** linear combination of the independent ones, we say the vector space is **n-dimensional**, and the superscript of the symbol $\mathcal{V}^{\mathbf{n}}$ indicates this. Recall that we have earlier mentioned this in terms of unique attributes of the objects. Hence, while representing each object as a vector we spell out its uniqueness in terms of (the same) n number of "components"; neither less, nor more. Needless to say, the components must be such that the set of vectors does not loose its independence; v_a, for instance, differs from every other vector is readily seen from the 25 zeros in the vector.

Example 3 [contd...]:

Let us now look at the word "coffee." Since it has a 'c', two 'e's, two 'f's, and a 'o', we write it as

$$v_c + 2v_e + 2v_f + v_o \ \in \ \mathcal{A}^{26}$$

Of course, we need to process further to make the new vector in non-negative integers look like the english word, but we do not go that far as of now.

As a next step, we make a *matrix* out of the *l.i.* independent set of n vectors but simply placing them as columns as follows:

$$\mathcal{A} \ = \ \begin{bmatrix} v_a & v_b & \cdots & v_y & v_z \end{bmatrix}$$

An immediate observation is that this is a diagonal matrix of size 26×26, with the alphabet along the principal diagonal. We use this matrix to premultiply the vector:

$$\begin{bmatrix} 0 & 0 & 1 & 0 & 2 & 2 & \cdots & 1 & \cdots & 0 \end{bmatrix}^T$$

where each row corresponds to a letter of the alphabet, typically arranged in the alphabetical order. In other words, for a 'c' we have a 1 in the 3^{rd} row, a 2 in the 5^{th} row, and so on to represent the word 'coffee' as $v_c + 2v_e + 2v_f + v_o$ with 4 non-zero entries in the vector.

Remark: The vector space has english words, including the alphabet, and the field has non-negative integers. Thus, technically speaking, every vector in $\mathcal{V}^n$ is mapped to another vector whose n-components are drawn from the field $\mathcal{F}$.

Following this shall be an important question the reader should have raised – does every vector space contain a unique set of $l.i.$ vectors? Since every vector space, e.g., english dictionary as above, has a large number of distinct vectors each of which must have been made up of elementary vectors there always exist at least one set of $l.i.$ elementary vectors in $\Re^n$:

$$e_i = \begin{cases} 1 & \text{in the } i^{th} \text{ row} \\ 0 & \text{rest of the rows} \end{cases} \tag{29}$$

The number of such vectors is also identical to the dimension n of the space, meaning that rest of the vectors in the space may be generated as linear combinations of the elementary vectors. Going back to example 3, however, it is interesting to note that one of the elementary vectors, e.g., v_a may be replaced by a linear combination of v_a and some other vectors, say, v_b and v_c, such that the new set of n vectors is linearly independant. For instance, we may easily verify using the definition, that the set

$$\{v_a,\, v_b,\, v_c,\, v_d,\, v_{ef},\, v_f,\, \cdots,\, v_z\} \quad \text{with } v_{ef} = v_e + v_f$$

is also $l.i.$

Hence, there could be infinitely many sets of $l.i.$ vectors. Moreover every such set may be represented by a square matrix, though not necessarily (why?) diagonal. Accordingly, other vectors have different representations – coffee $= v_c + 2v_{ef} + v_o$ with only three non-zero entries.

This observation from the above example leads to a very important concept. Any set of n linearly independent vectors of a n-dimensional vector space $\mathcal{V}^n$ constitutes a **basis** $\mathcal{B}$. Every other vector, a linear combination of this set of vectors, has a unique representation in $\mathcal{F}^n$ with reference to the basis[5] in context.

[5] Courtesy: Anu Garg, wordsmith.org.
What can one do with just seven letters? You'd think that would be limiting, but music made with those seven notes A-G can move the world. It's not the number of notes on your keyboard or the number of letters in your alphabet. It's what you do with them, how you

Definition: Isomorphism

An isomorphism φ from a vector space $\mathcal{V}$ to another vector space $\mathcal{V}'$, both over the same field $\mathcal{F}$, is a bijective[6] map $\varphi : \mathcal{V} \to \mathcal{V}'$ compatible with the operations $+$ and $\cdot$, i.e.,

$$\varphi \,:\, \varphi\left(\sum_i \alpha_i v_i\right) = \sum_i \left(\alpha_i \varphi\left(v_i\right)\right) \tag{30}$$

A very simple example is to view the set $\mathcal{C}$ of complex numbers as a real vector space $\Re^2$, i.e., the map $\varphi : \Re^2 \to \mathcal{C}$ sending an ordered pair $(r_1,\, r_2)$ of real numbers to $r_1 + jr_2$ is an isomorphism. Electrical circuit theory is completely based on this isomorphism.

Along similar lines, the space $\mathcal{F}^n$ of n-dimensional vectors drawn from the field $\mathcal{F}$ is isomorphic to the space $\mathcal{V}^n$.

It follows that it is mandatory to mention the basis whenever we address a vector; needless to say, there are infinitely many bases (plural for basis) to choose from and other vectors put on different appearances. Perhaps the best example to cite here is the number 10. By default, we think of decimal number system and it is *ten* for us, but for a machine, which works in binary number system by its default, it is *two*.

If $\mathcal{V}^n$ is a vector space over a field $\mathcal{F}$, and we chose $\mathcal{B}$ as the basis, then every vector $v_i \in \mathcal{V}^n$ is formally written as $\mathcal{B}f_i$ where $f_i \in \mathcal{F}^n$.

Thus, the vector space is very intimately associated with the underlying field. For instance, in the aforementioned x-y plane example, $\Re^2$ denotes 2-dimensional vector space built using the cartesian product of real numbers taken from the field $\mathcal{F} = \Re$, and hence we write $\Re^2 = \Re \times \Re$.

If we consider an arbitrary vector $v \in \mathcal{V}^n$, and two bases $\mathcal{B}_1$ and $\mathcal{B}_2$, then the vector's representations $\mathcal{B}_1 f_1$ and $\mathcal{B}_2 f_2$ respectively. Since it is the same vector we are referring to, therefore,

$$\mathcal{B}_1 f_1 = \mathcal{B}_2 f_2$$

and hence, for instance,

$$f_1 = \mathcal{B}_1^{-1} \mathcal{B}_2 f_2 \tag{31}$$

arrange them, that counts... Consider the alphabet of life. All life on Earth in its almost infinite variety of species and individual organisms is made up of just four letters: A, C, G, and T (the nucleic acids adenine, cytosine, guanine, and thymine), which, arranged in an endless number of sequences, make up DNA.

[6] If the map φ is not assumed to be a bijection it becomes a homomorphism.

Notice that the inverse exists since any basis is a $l.i.$ set.

A funny example here could be looking at, say a column of 4 letters in ascending order (basis $\mathcal{B}_1$) and descending order (basis $\mathcal{B}_2$):

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} A \\ M \\ P \\ Z \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} Z \\ P \\ M \\ A \end{bmatrix}$$

While we have a natural basis, such as the x- and y- axes, that spans every vector space, most of the time some other basis comes handy in computations and interpretations, and we need to welcome change. We will next establish that this is going to happen because each possible basis $\mathcal{B}_i$ possesses a certain *characteristic* feature, which is in turn is due the fact that the same vector v appears differently in different bases. We choose different bases depending on the context, and we may perform a **change of basis** operation to suit our requirement.

Since any basis is a square matrix of size $n \times n$, we will next explore the characterization of square matrices.

1.2 The Characteristic of a Square Matrix

Suppose we take a pair of rows[7] of a matrix A:

$$\begin{bmatrix} a_{i1} & a_{i2} & \cdots & a_{in} \end{bmatrix} \text{ and } \begin{bmatrix} a_{(i+1)1} & a_{(i+1)2} & \cdots & a_{i(i+1)n} \end{bmatrix}$$

where the former row is assumed to have all non-zero elements. We may notice something interesting here – if anyone of the elements of the $(i+1)^{th}$ row, $a_{(i+1)j} = 0$ then the two rows may be proven to be linearly independent of each other, i.e., $\alpha_1 r_1 + \alpha_2 r_2 = 0 \leftrightarrow \alpha_1 = \alpha_2 = 0$. By rearranging, i.e., by permuting[8], we may bring the zero of the latter row to the leading position, making $a_{(i+1)1} = 0$.

Extending this idea, we realize that all the rows of a diagonal matrix are necessarily independent, provided none of the diagonal elements is zero, and it is more practical to replace the diagonal matrix with a triangular matrix; this does not require any computation, it is by mere inspection. In other words, given n rows we typically leave the first row untouched and perform the "scale-and-add" operations on, i.e., linear combinations of, rows

[7] All the arguments presented here may be made in terms of columns of the matrix as well.

[8] what I call item #1, you may call it item #7, someone else may call it item #5 and so on

2 to n in such a way that for any k^{th} row $(k = 2, 3, \cdots, n)$, all the elements $a_{k1} \cdots a_{k(k-1)}$ are progressively made zero and check if the pivot a_{kk} does not become zero; elements $a_{k(k+1)} \cdots a_{kn}$ can be arbitrary. This was intuitively arrived at in eqn. (3). The following example 4^{th} order matrix illustrates the idea.

Example 4: Let the storekeeper has 4 items on sale, and the leader of buyers, let us call her buyer #1, devises the strategy – each buyer $\#i$ would go to the store, pretends that she does not know the others, and buys a_{ij} units of item $\#j$ according to the matrix

$$A = \begin{bmatrix} 4 & 4 & 3 & 4 \\ 4 & 3 & 4 & 2 \\ 3 & 4 & 4 & 5 \\ 1 & 1 & 1 & 2 \end{bmatrix}$$

and back home the leader would perform the following scale-and-add operations:

SA1. $r_2 = r_2 - r_1$ to make a_{21} zero,

SA2. $r_3 = 4 \cdot r_3 - 3 \cdot r_1$ to make a_{31} zero, followed by $r_3 = r_3 + 4 \cdot r_2$ to make a_{32} zero.

SA3. $r_4 = 4 \cdot r_4 - r_1$ to make a_{41}, a_{42} zero, followed by $r_4 = 11 \cdot r_4 - r_3$ to make a_{43} zero.

to get

$$A = \begin{bmatrix} 4 & 4 & 3 & 4 \\ 4 & 3 & 4 & 2 \\ 3 & 4 & 4 & 5 \\ 1 & 1 & 1 & 2 \end{bmatrix} \xrightarrow{SA1} \begin{bmatrix} 4 & 4 & 3 & 4 \\ 0 & -1 & 1 & -2 \\ 3 & 4 & 4 & 5 \\ 1 & 1 & 1 & 2 \end{bmatrix}$$

$$\xrightarrow{SA2} \begin{bmatrix} 4 & 4 & 3 & 4 \\ 0 & -1 & 1 & -2 \\ 0 & 0 & 11 & 0 \\ 1 & 1 & 1 & 2 \end{bmatrix} \xrightarrow{SA3} \begin{bmatrix} 4 & 4 & 3 & 4 \\ 0 & -1 & 1 & -2 \\ 0 & 0 & 11 & 0 \\ 0 & 0 & 0 & 44 \end{bmatrix}$$

a triangular matrix eventually, and the prices of the items are uniquely determined by back substitution.

Nevertheless, it is also quite possible that the storekeeper being an intelligent guy senses the idea and plays a small trick – he convinces the first buyer and sells her an additional unit of item 3; likewise, he sells one less unit of item 4 to buyer #3, and declares that item 3 is out of stock for buyer #4; purposefully, he avoids disturbing the diagonal transactions a_{ii}. Needless to say, the honest storekeeper billed the buyers exactly. Once the team goes back they find, fallen neatly into the trap, the matrix:

$$A = \begin{bmatrix} 4 & 4 & 4 & 4 \\ 4 & 3 & 4 & 2 \\ 3 & 4 & 4 & 4 \\ 1 & 1 & 0 & 2 \end{bmatrix}$$

does not help them. Why? One should quickly work out and verify that the last row vanishes in an attempt to make a_{43} zero. More formally we see the linear combination

$$\sum \alpha_i r_i = 2 \cdot r_1 - r_2 - r_3 - r_4 = \begin{bmatrix} 0 & 0 & 0 & 0 \end{bmatrix}$$

with none of the α_i equal to 0.

Thus, the pivots are *crucial* to the solution: in any case they must not be removed (altogether) from the matrix if we are heading towards a unique solution. In fact, the diagonal elements which actually help us establish linear independence when a triangular form is sought, are more likely to be vulnerable, during the scale-and-add operations, as suggested in the previous example. This calls for a sensitivity analysis by asking – is there something that, when removed (either directly or through the scale-and-add operations) from all the diagonal elements, makes the rows dependent, and hence the determinant zero, leading to an explosion of solutions? We rephrase this question as a meaningful equation:

$$?\exists \lambda \ : \ \Delta_n = |(A - \lambda I)| = 0 \tag{32}$$

A straight answer may be obtained if we have the diagonal (or triangular) matrix in our mind, in which case we get the following polynomial equation

$$\begin{aligned} \Delta_n &= (a_{11} - \lambda) M_{11} \\ &= (a_{11} - \lambda)(a_{22} - \lambda) M_{22} \\ &\vdots \\ &= \Pi_{i=1}^{n} (a_{ii} - \lambda) = 0 \end{aligned} \tag{33}$$

Thus, it is not just one, but n possible complex scalars, anyone of them when subtracted from all the diagonal elements would render the rows dependent. Here M_{ii} are called the **principal minors**[9].

[9] These are cofactors as well.

In any case - whether the square matrix is arbitrary, or triangular, or simply diagonal, the outcome is:

$$\chi_A = \lambda^n - \gamma_{n-1}\lambda^{n-1} + \gamma_{n-2}\lambda^{n-2} + \cdots + (-1)^n\gamma_0 = 0 \tag{34}$$

called the **characteristic equation**. Assuming that the matrices of our interest are real, i.e., the rows/columns of the matrices belong to the vector space $\Re^n$ and hence the matrices belong to $\Re^{n\times n}$, the coefficients γ_i of the polynomial would be all real, and hence the n roots λ_i, called by various names – characteristic values, proper values, or **eigenvalues**[10] – are, in general, complex numbers. Let us quickly observe that the vectors over the field $\mathcal{F} = R$ are n-tuples $\Re^n$, followed by the matrices belonging to $\Re^{n\times n}$; even the determinant is a map $\Re^{n\times n} \to \Re$, but not all results of composite operations need to belong the same field. Hence, we do consider *extended fields*, for instance the set of complex numbers $\mathcal{C}$ in this case, which include the real numbers; we call $\mathcal{C}$ as the algebraic closure of the real field $\Re$. Thus, the set of eigenvalues is the characteristic of a square matrix in the sense that it contains important information about the nature of the matrix.

Looking at the eigenvalues from this perspective, they are the n critical values the removal of any one of which from all the diagonal elements may not permit the matrix A to perform its intended transformative operation, here the change of basis. And, that must be the reason behind choosing the German word. In a general matrix, unlike a diagonoal or triangular one, no eigenvalue is explicitly visible and hence it is *latent* and needs to be extracted as a root of the characteristic polynomial. Lastly, we later on witness that any function of the square matrix A depends on the eigenvalues rather than on the elements of the matrix, and in that sense the word *characteristic* appears to be more appropriate. Understanding this way, what else could be the mechanism, other than setting $|A - \lambda I| = 0$, to bring out the eigenvalues?

Example 5: Let

$$A = \begin{bmatrix} 0 & 1 \\ -2 & 3 \end{bmatrix} \in \Re^{2\times 2}$$

The characteristic equation and the eigenvalues are

$$\lambda^2 - 3\lambda + 2 \;:\; \lambda_{1,2} = 1,\ 2$$

Example 6: Let

$$A = \begin{bmatrix} 3 & 0 \\ 4 & -1 \end{bmatrix} \in \Re^{2\times 2}$$

[10] Almost every combination of the adjectives proper, latent, characteristic, eigen and secular, with the nouns root, number and value, has been used in the literature for what we call a proper value – Paul R. Halmos [9, p. 102].

The characteristic equation and the eigenvalues are

$$\lambda^2 - 2\lambda - 3 \,:\, \lambda_{1,2} = \{-1,\, 3\}$$

Since the characteristic polynomial $\chi_A(\lambda)$ is the determinant of $A - \lambda I$, it readily follows that

$$|A| = \chi_A(0) = (-1)^n \gamma_0$$

i.e., the determinant of any matrix is the product of its eigenvalues. Therefore, the rows are all independent if none of the eigenvalues is zero, and this leads to another interesting connection – the rank of any square matrix, e.g., basis of a n-dimensional space, is the number of non-zero eigenvalues. While it is readily visible in the case of diagonal and triangular matrices, it is a fact that the rank of $A - \lambda_i I$ looses its rank depending on the geometric multiplicity of the eigenvalues; for instance, the identity matrix (the natural/standard basis) has $\lambda = 1$ repeated n times and the rank of $A - \lambda I = 0$. In general, if the rank of a matrix whose columns are drawn from a vector space of dimension n, is $n - c$, we call c as the *nullity* of the matrix, and we have one of the fundamental theorems, the rank theorem: rank + nullity = dimension n. It is indeed a matter of serious concern when the characteristic polynomial as multiple roots, but in the interest of the big picture in this brief appendix we do not bring it here.

Furthermore, i is easy to verify that, if λ_i are the eigenvalues of A, the reciprocals $1/\lambda_i$ would be the eigenvalues of A^{-1}.

There is much more to the eigenvalues. However, we first need to understand a little more formally about the transformations, either among the vectors of the same space (e.g., change of basis) or between vectors of two different spaces (e.g.,). Soon after this we resume our discussion on eigenvalues.

1.3 Linear Transformations

Going by our theme puzzle, we have on one side a vector of bills paid by the students, and on the other side is a vector of unit price of each of the n items. Indeed, what we presumed and further developed is that these two vectors, and for that matter any pair of vectors one from each category, are linearly related, i.e., every component of one of the vectors is a linear combination of the all the components of the other vector. In other words, we *transformed* one vector into another.

More formally, if there are two spaces $\mathcal{U}^m$ and $\mathcal{V}^n$, a linear transformation $\mathbb{L}$ is a map

$$\mathbb{L} : u \in \mathcal{U}^m \to v \in \mathcal{V}^n \tag{35}$$

satisfying the following properties:

$$\mathbb{L}(\bar{0}) = \bar{0} \quad \text{and} \quad \mathbb{L}\left(\sum_i \alpha_i u_i\right) = \sum_i \alpha_i \mathbb{L}\left(u_i\right) \tag{36}$$

It is not difficult to see that $\mathbb{L}$ readily takes the form of a matrix of size $n \times m$, pre-multiplying the vector $u \in \mathcal{U}^m$ to give us $v \in \mathcal{V}^n$. Typically, this matrix is called as the matrix of $\mathbb{L}$ with respect to the bases of $\mathcal{U}$ and $\mathcal{V}$. As we elaborate more in the following pages, different choices of the bases lead to different matrices.

Example 7: A polynomial of degree n may be represented as a vector in the the standard basis as

$$p_n(x) = I_n \begin{bmatrix} p_n \\ p_{n-1} \\ \vdots \\ p_0 \end{bmatrix} \in \mathcal{P}^{n+1}$$

where the columns of the basis I_n, from left to right, indicate $x^n, x^{n-1}, \cdots, x^0$ and the vector consists of all the coefficients p_i in the decreasing order of power of x_i. Thus, this is a $(n+1)$-dimensional space. By differentiating the polynomial (w.r.t. x) once, we define a linear transformation from $\mathcal{P}^{n+1}$ to $\mathcal{P}^n$.

With reference to the properties of linear transfomation, we have the following mportant definitions.

Definition: Sub-space

Let $\mathcal{V}^n(+,\cdot)$ be a vector space over $\mathcal{F}$ and let $\mathcal{W} = \{w_1, w_2, \cdots\}$ be a nonempty subset of $\mathcal{V}$. Then,

$$\mathcal{W}(+,\cdot) \text{ is a subspace of } \mathcal{V} \Leftrightarrow \forall \alpha_i \in \mathcal{F} \sum_i \alpha_i w_i \in \mathcal{W} \tag{37}$$

Remarks:

1. If $\mathcal{W}$ is a subspace of $\mathcal{V}$, then $\mathcal{W}$ contains the zero vector $\bar{0}$ of $\mathcal{V}$.

2. The dimension of $\mathcal{W} \leq n$. $\mathcal{V}$ is clearly a subspace of itself. The set $\{\bar{0}\}$, consisting of only the zero vector, is also a subspace of $\mathcal{V}$, called the zero subspace. These are called the trivial subspaces of $\mathcal{V}$; all others are generally called as proper subspaces.

Definition: Kernel

The **kernel** of a linear transformation is

$$ker(\mathbb{L}) = \{u \in \mathcal{U} \, : \, \mathbb{L}(u) = \bar{0}\} \tag{38}$$

Notice that if $\mathbb{L}$ is a matrix of size $m \times n$, for instance with $m < n$, then there would be non-trivial solutions v_i, in addition to the trivial solution $\bar{0}$, that satisfy $\mathbb{L}(v) = \bar{0}$

Definition: Image

The **image** of a linear transformation is

$$im(\mathbb{L}) = \{v \in \mathcal{V} \, : \, \mathbb{L}(u) = v \text{ for some } u \in \mathcal{U}\} \tag{39}$$

Example 8: Let us consider $\mathcal{U}^3$ and $\mathcal{V}^2$ over the field $\Re$ with

$$\mathbb{L} = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 9 \end{bmatrix}$$

then

$$ker(\mathbb{L}) = \text{span of the vectors } [1 \;\; 1 \; -1]^T \in \mathcal{U}^3$$

i.e., $ker(\mathbb{L})$ is a sub-space of $\mathcal{U}$ – only those vectors that are scaled versions of $[1 \;\; 1 \; -1]^T$, and hence its dimension is 1.

For $im\mathbb{L}$, we see that the span of $\mathbb{L}\,(u \in \mathcal{U}^3)$ is a subspace of $\mathcal{V}$ of dimension 2 – the basis of these image vectors is restricted by the linear combinations induced by $\mathbb{L}$ and hence not all vectors of $\mathcal{V}$ are generated by this transformation for any $u \in \mathcal{U}$.

Thus, it is customary to say that the set $\mathcal{L}$ of linear transformations is a vector space of dimension $m \times n$ over the field $\mathcal{F} = \mathcal{F}_U \cup \mathcal{F}_V$. Moreover,

$$\dim\left(ker(\mathbb{L})\right) + \dim\left(im(\mathbb{L})\right) = \dim\left(\mathcal{U}^m\right) \tag{40}$$

This is the formal version of the aforementioned rank theorem *nullity + rank = dimension*.

Earlier we mentioned that the matrix $\mathbb{L}$ is with respect to the bases of $\mathcal{U}$ and $\mathcal{V}$. If we study changes of basis in the respective spaces, we will reap several benefits. We will basically ask: What happens to the matrix $\mathbb{L}$ of the transformation if we make other choices of bases. If $\mathcal{B}_U$ and $\mathcal{B}_V$ are the new bases such that

$$\mathcal{B}_U u' = u \;\text{ and }\; \mathcal{B}_V v' = v$$

where u and v are originally seen from their standard basis, then a linear transformation between the two spaces results in

$$\mathcal{B}_V v' = \mathbb{L}\mathcal{B}_U u'$$

If $\mathbb{L}'$ is the matrix with respect to the new bases, then

$$\mathbb{L}' = \mathcal{B}_V^{-1}\,\mathbb{L}\,\mathcal{B}_U \;\text{ such that }\; v' = \mathbb{L}'u' \tag{41}$$

In fact, we may interpret $\mathcal{B}_V^{-1}\mathbb{L}\mathcal{B}_U$ as the matrix obtained from $\mathbb{L}$ by a succession of row and column operations, e.g., leading to a matrix that has a lot of zeros, like a diagonal matrix, so that multiplication by such matrices is easy to describe. Therefore, it is but natural to hunt for bases such that $\mathbb{L}'$ becomes significantly simplified. We will provide an example after a few lines. One would also, now, observe the power of working in structures like vector spaces without fixed bases.

Next, if a linear transformation maps a vector in $\mathcal{V}$ to itself, i.e., $\mathbb{L} : \mathcal{V} \to \mathcal{V}$, then it is referred to as a linear *operator* on $\mathcal{V}$. Obviously, in the matrix representation, it is a square matrix. And, as a pre-multiplying matrix, the basis matrix is a linear operator on the space $\mathcal{F}^n$ of column vectors. Having learnt about transformations between two arbitrary spaces, we now quickly adapt the same to linear operators. We need only one basis $\mathcal{B}$ for $\mathcal{V}$, and use it in place of both the bases $\mathcal{B}_U$ and $\mathcal{B}_V$. Thus, if

$$v_2 = \mathbb{L}v_1,\; v_1, v_2 \in \mathcal{V}^n$$

and we choose a new basis $\mathcal{B}$, with respect to this new basis the operator becomes

$$\mathbb{L}' = \mathcal{B}^{-1}\,\mathbb{L}\,\mathcal{B} \tag{42}$$

What we technically call as *change of basis*, is a shift in our point-of-view to look at the unknowns differently, but plausibly in a more computation-friendly way.

In general, we say that a square matrix A is **similar** to A' if $A' = M^{-1}AM$ for some non-singular M. Once again, it is natural to hunt for a similar matrix which is particularly simple. However, here we face certain restrictions since we have only one basis, and therefore one matrix to work with. Nevertheless we proceed in the following manner.

Recall that subtracting an eigenvalue λ_i from the diagonal elements makes the matrix rank deficient. In other words, looking at the rows, there exists a set of scalars in $\mathcal{F}$ such that the linear combination of all the n rows is zero, and we may quickly observe that for each eigenvalue λ_i,

$$\exists \mathbf{l_i} \neq 0 \in \mathcal{C}^n \;:\; \mathbf{l_i}^T \left[A - \lambda_i I\right] = 0 \;\; i = 1, 2, \cdots, n \tag{43}$$

Such *vectors*, n in number and independent of each other, are called **left eigenvectors corresponding to the eigenvalues**. Basically, each $\mathbf{l_i}$ is a stack of all those scalars such that the linear combination of the rows of $A - \lambda_i I$ is zero.

Likewise, we may also observe that for each eigenvalue λ_i

$$\exists \mathbf{r_i} \neq 0 \in \mathcal{C}^n \ : \ [A - \lambda_i I]\, \mathbf{r_i} = 0 \quad i = 1, 2, \cdots, n \tag{44}$$

or, equivalently

$$A\mathbf{r_i} = \lambda_i \mathbf{r_i}$$

Each of these n vectors $\mathbf{r_i}$ are called **right eigenvectors**, conventionally just the **eigenvectors corresponding to the eigenvalues**.

Example 6 contd...: For

$$A = \begin{bmatrix} 3 & 0 \\ 4 & -1 \end{bmatrix} \in \Re^{2\times 2}$$

The eigenvalues of this lower triangular matrix were found to be

$$\lambda_{1,2} = \{-1,\ 3\} \ \text{ along the diagonal}$$

Corresponding to the eigenvalue $\lambda_1 = -1$,

$$A - \lambda_1 I = \begin{bmatrix} 4 & 0 \\ 4 & 0 \end{bmatrix}$$

A quick observation again – rank of $A - \lambda_1 I$ is 1. And, we may compute

$$\mathbf{l_1} = \begin{bmatrix} 1 \\ -1 \end{bmatrix} \ \text{ and } \ \mathbf{r_1} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

so that

$$\mathbf{l_1}^T A = \lambda_1 \mathbf{l_1}^T \ \text{ and } \ A\mathbf{r_1} = \lambda_1 \mathbf{r_1}$$

Likewise, for $\lambda_2 = 3$

$$\mathbf{l_2} = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \ \text{ and } \ \mathbf{r_2} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$$

are the left- and right-eigenvectors respectively.

In the general $n \times n$ case, if we compute all the n right eigenvectors, and if we are lucky to see the set to be *l.i.*, we may proceed to compose the matrix:

$$M = [\mathbf{r_1} \ : \ \mathbf{r_2} \ : \ \cdots \ : \ \mathbf{r_n}] \tag{45}$$

known as the **eigenbasis** of A. Consequently, we may write $A\mathbf{r_i} = \lambda_i \mathbf{r_i}$ as

$$
\begin{aligned}
AM &= [\lambda_1 \mathbf{r_1} \,:\, \lambda_2 \mathbf{r_2} \,:\, \cdots \,:\, \lambda_n \mathbf{r_n}] && (46)\\
&= M \begin{bmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 & \cdots & 0 \\ \vdots & & \ddots & \\ 0 & 0 & \cdots & \lambda_n \end{bmatrix} && (47)\\
&= M\Lambda && (48)
\end{aligned}
$$

where Λ denotes the diagonal matrix as above. Rearranging terms, we have

$$M^{-1}AM = \Lambda \tag{49}$$

i.e., any square matrix, particularly the one having distinct eigenvalues, is *similar* to a diagonal matrix; notice that the diagonal matrix is the simplest one with a lot of zeros.

Thus, if A is a linear operator and v_1 and v_2 are now seen from the basis $\mathcal{B} = M^{-1}$, then the operation is most simple as it amounts to scaling each of the components of $\mathcal{B}v_1$ by λ_i to get the corresponding component of $\mathcal{B}v_2$.

While figuring out the independent vectors $\mathbf{r_i}$ and composing a non-singular M are indeed possible, it is not so easy in practice since inverting a matrix tends to be prohibitively computationally expensive – it needs $n!$ multiplications to compute the determinant and $n \times n!$ multiplications to compute the n^2 cofactors resulting in a whopping $(n+1)!$ multiplications, using Laplace's expansion. Therefore, orthonomal matrices computed using, for instance, successive QR-decompositions[11] on A, are preferred, and the process involves a sequence of updates starting from A.

Another noteworthy point here is that given a matrix A, one may attempt and work out to verify that pre-multiplying A by a matrix amounts to performing row operations. One may also verify that post-multiplying A by the inverse (of the earlier matrix) amounts to performing the dual set of operations on the columns. As a consequence, all the matrices $A, A_1 \cdots A_i$, and the diagonal matrix Λ (if we are indeed able to arrive at), have identical set of eigenvalues λ_i. This is very easy to establish – if two matrices A_1 and A_2 are similar in the sense that

$$\text{if} \;\; M^{-1}A_1M = A_2$$

[11]More about this in a later section.

for some non-singular M, then

$$\begin{aligned} \chi_{A_1} = 0 \;\; &= \;\; |\,(A_1 - \lambda I)\,| \\ &= \;\; |\,\left(MA_2M^{-1} - \lambda MIM^{-1}\right)| \\ &= \;\; |M\,(A_2 - \lambda I)\,M^{-1}| \\ &\Rightarrow \;\; |\,(A_2 - \lambda I)\,| \; = \; 0 \\ &= \;\; \chi_{A_2} \end{aligned}$$

In other words, similar matrices have identical characteristic equations, i.e., same set of eigenvalues (and hence the same rank).

The computation of eigenvectors corresponding to the eigenvalues such that

$$A\mathbf{r_i} = \lambda_i \mathbf{r_i}, \quad i = 1, 2, \cdots, n$$

is not a trivial task; in fact, in any classroom a very small number of students would be able to solve the above equation for $n = 3$. While $\mathbf{r_i} = 0$ is the trivial solution, we need at least one non-trivial solution. By and large the difficulty arises because we attempt to solve a rank-deficit system of equations, and interestingly, every student is baffled at the same equation repeated twice or more. In addition to having a routine, an insight can be obtained by looking at the upper triangular matrix, with the eigenvalues along the principal diagonal:

$$A = \begin{bmatrix} \lambda_1 & a_{12} & \cdots & a_{1n} \\ 0 & \lambda_2 & & a_{2n} \\ \vdots & & \ddots & \\ 0 & 0 & \cdots & \lambda_n \end{bmatrix}$$

With $\lambda = \lambda_1$, the first column of $A - \lambda I$ vanishes and hence the only possible non-trivial eigenvector is

$$\mathbf{r'_1} = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

Extending this idea to the rest of the eigenvalues, we get

$$M' = \begin{bmatrix} 1 & * & \cdots & * \\ 0 & 1 & \cdots & * \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{bmatrix}$$

where $*$ denotes a possibly non-zero element making M upper triangular as well.

Thus, computationally we fix $\mathbf{r'_{ii}} = 1$ and solve linear equations of size $n-1$ for each eigenvalue λ_i.

Once this set of eigenvectors is obtained, given the similarity, the required eigenvectors of A may be obtained by the transformation:

$$\mathbf{r_i} = Q\mathbf{r'_i},\ i = 1, 2, \cdots, n$$

Let us now go back to our group of buyers and look at the original linear system of equations $Ax = b$ from the above perspective; for the sake of convenience we assume A has full rank, having successfully escaped from the shopkeeper's tricks. An important remark needs to be made here. While the strategic shopping was going on, back home the leader had been busily scaling, adding, and rearranging the rows of the matrix A, as well as those of the column b – essentially pre-multiplying the matrix with another – to arrive at an upper triangular matrix that has no zeros along the diagonal; columns were consciously left untouched. In such a case, the eigenvalues of the triangular matrix could be different from those of A. The goal has always been to obtain the unique solution x of the linear system. More details on this line of argument may be found in a later section on QR-decomposition.

Moreover, if the group made column- as well as row-operations on A such that $M^{-1}AM = A'$ is upper-triangular, it is equivalent to first transforming the unknown vector x to another vector Mz and the vector b to Mb' on the right hand side. Effectively, we get

$$A'z = b'$$

with identical set of eigenvalues for A and A'. Thus by change of basis, we pretty well know that solving the linear system by back substitution is straight - $\mathcal{O}\left(n^2\right)$ multiplications; once we obtain z by back substitution, we transform back to obtain $x = Tz$ in another $\mathcal{O}\left(n^2\right)$ multiplications. Thus, including the cost of computing the triangular matrix, the total number of multiplications and additions of products (more generally, the total number of floating point operations or flops) is $\mathcal{O}\left(n^3\right)$, which is extremely small compared with $(n+1)!$ multiplications needed to compute the solution via matrix inversion. Moreover, when A has full rank, none of the diagonal elements of A' is a zero, i.e., A' is also a bijection and hence a unique z is guaranteed. To eulogize the similarity transformation once more, we will briefly look at the matrix exponential e^A that plays a key role in the solution of the state equation of continuous-time systems. This function of a square matrix can be defined as a natural extension of its scalar counterpart –

$$e^A = I + A + \frac{A^2}{2!} + \frac{A^3}{3!} + \cdots + \frac{A^i}{i!} \cdots$$

Keeping aside for the time being that this is an infinite series, one finds that this function basically depends on the eigenvalues of A directly. For instance, one may have a flash-thought:

$$\begin{aligned} M^{-1}e^A M &= I + M^{-1}AM + \cdots \\ &= I + \Lambda + \cdots \\ &= e^{M^{-1}AM} \end{aligned}$$

And, after a sumptous meal the desserts are

1. The determinant is the product of all eigenvalues, with rank understood as the number of non-zero eigenvalues.
2. Two matrices connected via $A' = M^{-1}AM$ are said to be **similar** $\Leftrightarrow$ they possess the same set of eigenvalues, and hence the same rank.
   (a) If A and A' are similar, their trace, basically the sum of diagonal elements but turns out to be the sum of eigenvalues, is also the same; however, the converse is not always true.

The matrices A and A' thus have a common means of formation through the eigenvalues, and their *similarity* under the eyepiece of eigenbasis T is uncanny. Every square matrix is a symphony with the spectrum of its eigenvalues as the underlying set of chords. To this end we will next see one of the most beautiful theorems in linear algebra.

1.4 The Cayley-Hamilton Theorem

Let us have a matrix A, and a transformation M such that $M^{-1}AM = \Lambda$. For each of the n eigenvalues we have a characteristic equation:

$$\lambda_i^n - \gamma_{n-1}\lambda_i^{n-1} + \gamma_{n-2}\lambda_i^{n-2} + \cdots + (-1)^n\gamma_0 = 0$$

These n equations can be packed into a matrix equation

$$\Lambda^n - \gamma_{n-1}\Lambda^{n-1} + \gamma_{n-2}\Lambda^{n-2} + \cdots + (-1)^n\gamma_0 I = [0] \tag{50}$$

where $[0]$ is the zero matrix. Using M,

$$\left(M\Lambda M^{-1}\right)^n - \gamma_{n-1}\left(M\Lambda M^{-1}\right)^{n-1} + \cdots + (-1)^n\gamma_0 MIM^{-1} = M[0]M^{-1}$$

which is none other than

$$A^n - \gamma_{n-1}A^{n-1} + \gamma_{n-2}A^{n-2} + \cdots + (-1)^n\gamma_0 I = [0] \tag{51}$$

This is the famous theorem stated as

Every square matrix satisfies its own characteristic equation.

Just for instance, we may *compute*

$$A^{-1} = -\frac{1}{(-1)^n\gamma_0}\left(A^{n-1} - \gamma_{n-1}A^{n-2} + \gamma_{n-2}A^{n-3} + \cdots + \gamma_1 I\right)$$

For that matter any function of a square matrix may be reasonably well computed as a linear combination of positive powers of the matrix, using this theorem.

Example 9: Let

$$A = \begin{bmatrix} 3 & 0 \\ 4 & 1 \end{bmatrix}$$

The characteristic equation is

$$\lambda^2 - 4\lambda + 3 = 0$$

and with little effort one can readily verify that

$$A^2 - 4A + 3I = [0]$$

from which

$$A^{-1} = -\frac{1}{3}(A - 4I) = -\frac{1}{3}\left(\begin{bmatrix} (3-4) & (0-0) \\ (4-0) & (1-4) \end{bmatrix}\right) = \frac{1}{3}\begin{bmatrix} 1 & 0 \\ -4 & 3 \end{bmatrix}$$

The Cayley-Hamilton theorem actually helps make infinite series such as e^A finite and allow an algorithmic computation.

We next look at a very important concept that indicates the *size*, in some sense, of a vector or a matrix in the space.

1.5 Vector and Matrix Norms

Here we will just provide the basic idea and encourage readers to get more details from the references. If a real number, a scalar, has an absolute value why not a vector or a matrix? Put it simply, is there a way for our buying team to get an idea of how expensive is the store?

A vector more generally is a mathematical object, say a matrix or a polynomial, and often we need to have a feel of how big or small is that in some sense of magnitude – how far away is the tip of the vector (visualized as in physics as an arrow in a geometric space of coordinates) from the origin? For instance a polynomial like $x^3 + 3x^2 + 4x + 7$ could be a generalization of, say 1347 in decimal system where $x = 10$; the question is to assess how big is such a number, for instance like the distance of $|x|$ from zero on the real line. As a generalization, we abstract this 4-digit number as a vector in a four-dimensional space and assign the coordinates 1–3–4–7. We may then use Pythogorean theorem and find the distance between this point and the origin as $\sqrt{1^2 + 3^2 + 4^2 + 7^2}$.

Thus, in general we define the **Euclidean Norm**, or more conveniently the 2-norm of a vector $x = [x_1 \ x_2 \ \cdots \ x_n]^T$ as

$$\|x\|_2 \triangleq \sqrt{\sum_{i=1}^{n} x_i^2} = \sqrt{x^* x} \tag{52}$$

where we assume $x \in \mathcal{C}$ and x^* is the complex-conjugate transpose of x. This expression is pretty close to that of statistical variance of a zero-mean data. And, this norm is frequently seen in optimal control problems.

More generally, we define a $p-$norm as follows.

$$displaystyle \|x\|_p \triangleq \left(\sum_{i=1}^{n} x_i^p \right)^{\frac{1}{p}} \tag{53}$$

Using the notion of the norm of a vector, we may define the norm of a matrix as follows; after all, a matrix can also be thought of a generalized number. A rather straight way to judge the size of a matrix is

$$\|A\| \triangleq \|Ax\| \ \text{ where } x \text{ is any arbitrary vector } : \|x\| = 1 \tag{54}$$

This is the size of the transformed vector, implicitly giving amplification factor of the matrix. It is interesting to see that if it is 2-norm of x, then

$$\|A\|_2 = \max_{\|x\|_2=1} \|Ax\|_2 \ = \ \sqrt{\sigma_{max}} \tag{55}$$

where σ_{max} is the largest eigenvalue of A^*A (also known as the singular value of A), where A^* is the complex-conjugate-transpose of A. As a quick extension one might also think of

$$\|A^{-1}\|_2 = \frac{1}{\max_{\|x\|_2=1} \|Ax\|_2} \ = \ \frac{1}{\sqrt{\sigma_{min}}} \tag{56}$$

where σ_{min} is the smallest sigular value of A. The spread of these singular values is usually considered in the singular value decomposition discussed in the next section.
We now define a vector of all 1s:

$$\mathbf{1} = [1 \;\; 1 \;\; \cdots \;\; 1]^T \tag{57}$$

using which we define two different norms that are practically more popular.

$$\|A\| \triangleq \text{largest component of } \; A\mathbf{1} \text{ or } \mathbf{1}^T A \tag{58}$$

The first one gives the largest among the sums of elements along the rows and the second one gives the largest column sum. At times, we define a matrix norm as:

$$\|A\| = \max\left\{A\mathbf{1},\, \mathbf{1}^T A\right\} \tag{59}$$

Earlier we mentioned about extended fields, and to accommodate operations on eigenvalues we need to have $\mathcal{C} \ni \Re$ as the underlying field for many practical matrices. Accordingly, a useful norm known as Frobenius norm is given by

$$\|A\|_F^2 \triangleq \sum_{i,j} |a_{ij}|^2 \; = \; \text{trace } \; (A^* A) \tag{60}$$

In general, design solutions resulting as matrices are not unique and physical considerations and constraints, such as actuator contraints or bandwidth constraints, are likely to demand a restriction on the size of the elements of these matrices, and these matrix norms come handy.

Any of these (vector as well as matrix) norms need to satisfy the following properties, for obvious reasons as well as an extension from the absolute value of scalars.

1. $\|x\| \nless 0$, $\|x\| = 0 \Leftrightarrow x = 0$
2. $\|\alpha x\| = |\alpha| \|x\|$
3. $\|x + y\| \leq \|x\| + \|y\|$, the classic triangle inequality; we can also have the backward triangle inequality, $\|x\| - \|y\| \leq \|x - y\|$.

1.6 The QR Decomposition

Let us consider a square matrix A of n columns, $\mathbf{a_1}, \cdots, \mathbf{a_n}$, treated as vectors in the standard basis $\mathbf{e_1}, \cdots, \mathbf{e_n}$. We wish to express each of these vectors in another orthonormal basis, $\mathbf{q_1}, \cdots, \mathbf{q_n}$. Each column $\mathbf{a_j}$ may then be expressed as the following linear combination:

$$\mathbf{a_j} = \sum_{i=1}^{n} r_{ij}\mathbf{q_i} \tag{61}$$

If the matrix Q is the composition of the new orthonormal basis, then

$$A = Q \begin{bmatrix} r_{11} & r_{12} & \cdots & r_{1n} \\ r_{21} & r_{22} & \cdots & r_{2n} \\ \vdots & & \ddots & \\ r_{n1} & & \cdots & r_{nn} \end{bmatrix} \tag{62}$$

where all the scalars r_{ij} make up the matrix R.

Now the problem is twofold: to construct the new orthonormal basis $\mathbf{q_1}, \cdots, \mathbf{q_n}$ as well as compute the scalars r_{ij}, given the vectors $\mathbf{a_1}, \cdots, \mathbf{a_n}$. At first, this would appear formidable. But, two simple and clever tricks would open the path. The first trick is to consider $r_{ij} = 0 \ \forall i > j$, making the R matrix an upper triangular one. This gives us the following column-wise progression:

$$\begin{aligned} \mathbf{a_1} &= r_{11}\mathbf{q_1} \\ \mathbf{a_2} &= r_{12}\mathbf{q_1} + r_{22}\mathbf{q_2} \\ \mathbf{a_3} &= r_{13}\mathbf{q_1} + r_{23}\mathbf{q_2} + r_{33}\mathbf{q_3} \\ &\vdots \\ \mathbf{a_n} &= \sum_{i=1}^{n} r_{in}\mathbf{q_i} \end{aligned} \tag{63}$$

Noting that $Q^{-1} = Q^T$, we may rewrite the above equation as

$$Q^T A = R$$

so that

$$r_{11} = <\mathbf{q_1}, \mathbf{a_1}> = \mathbf{q_1}^T\mathbf{a_1}$$

to begin with, and in general,

$$r_{ij} = <\mathbf{q_i}, \mathbf{a_j}> \tag{64}$$

Apparently, a way out to compute the coefficients, but the vectors appear to be still elusive. The mystery unravels beautifully with the master stroke:

$$displaystyle \text{Choose } \mathbf{q_1} = \frac{1}{\|a_1\|}\mathbf{a_1} \tag{65}$$

Easily, for $A \in \Re^{3\times 3}$ the algorithm follows:

$$
\begin{aligned}
r_{12} &= < \mathbf{q_1}, \mathbf{a_2} > \\
r_{22}\mathbf{q_2} &= \mathbf{a_2} - r_{12}\mathbf{q_1} \\
\mathbf{q_2} &= \frac{1}{r_{22} \; = \; \|a_2 - r_{12}q_1\|} r_{22}\mathbf{q_2} \\
r_{13} &= < \mathbf{q_1}, \mathbf{a_3} > \\
r_{23} &= < \mathbf{q_2}, \mathbf{a_3} > \\
r_{33}\mathbf{q_3} &= (\mathbf{a_3} - r_{13}\mathbf{q_1}) - r_{23}\mathbf{q_2} \\
\mathbf{q_3} &= \frac{r_{33}\mathbf{q_3}}{r_{33} \; = \; \| \, (\mathbf{a_3} - r_{13}\mathbf{q_1}) - r_{23}\mathbf{q_2}\|}
\end{aligned}
\qquad (66)
$$

Note that r_{13} could have been computed alongside r_{12}, and the vector a_3 could have been *partially* updated to $a_3 - r_{13}q_1$, and later completely updated to $r_{33}q_3$; the parenthesis in the last couple of lines elicit us write the following tiny code:

---

for $i = 1 : n$ {

$$r_{ii} = \|\mathbf{a_i}\|$$

$$\mathbf{q_i} = \frac{1}{r_{ii}} \cdot \mathbf{a_{ii}}$$

for $j = (i+1) : n$ {

$$r_{ij} = \mathbf{q_i^T} a_j$$

$$\mathbf{a_j} = \mathbf{a_j} - r_{ij}\mathbf{q_i}$$

}

}

Example 10: Let

$$A = \begin{bmatrix} 2 & 2 & 1 \\ 2 & 1 & 2 \\ 1 & 2 & 2 \end{bmatrix}$$

<u>$i = 1:$</u>

$$
\begin{aligned}
r_{11} &= \|\mathbf{a_1}\| = 3 \\
\mathbf{q_1} &= \frac{1}{3}\begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix}
\end{aligned}
$$

$$
\begin{aligned}
j = 2: \quad r_{12} &= \; < \mathbf{q_1}, \mathbf{a_2} > = \frac{8}{3} \\
\mathbf{a_2} &= r_{22}\mathbf{q_2} \; = \; \mathbf{a_2} - r_{12}\mathbf{q_1} \; = \; \frac{1}{9}\begin{bmatrix} 2 \\ -7 \\ 10 \end{bmatrix} \\
j = 3: \quad r_{13} &= \; < \mathbf{q_1}, \mathbf{a_3} > = \frac{8}{3} \\
\mathbf{a_3} &= (\mathbf{a_3} - r_{13}\mathbf{q_1}) \; = \; \frac{1}{9}\begin{bmatrix} -7 \\ 2 \\ 10 \end{bmatrix} \\
\underline{i = 2:} \qquad r_{22} &= \|\mathbf{a_2} - r_{12}\mathbf{q_1}\| \; = \; \frac{\sqrt{153}}{9} \\
\mathbf{q_2} &= \frac{1}{r_{22}} r_{22}\mathbf{q_2} \; = \; \frac{1}{\sqrt{153}}\begin{bmatrix} 2 \\ -7 \\ 10 \end{bmatrix} \\
j = 3: \quad r_{23} &= \; < \mathbf{q_2}, \mathbf{a_3} > = \frac{8}{\sqrt{153}} \\
\mathbf{a_3} &= r_{33}\mathbf{q_3} \; = \; (\mathbf{a_3} - r_{13}\mathbf{q_1}) - r_{23}\mathbf{q_2} \; = \; \frac{1}{153}\begin{bmatrix} -135 \\ 90 \\ 90 \end{bmatrix} \\
\underline{i = 3:} \qquad r_{33} &= \|(\mathbf{a_3} - r_{13}\mathbf{q_1}) - r_{23}\mathbf{q_2}\| \; = \; \frac{\sqrt{34425}}{153} \\
\mathbf{q_3} &= \frac{1}{r_{33}} r_{33}\mathbf{q_3} \; = \; \frac{1}{\sqrt{34425}}\begin{bmatrix} -135 \\ 90 \\ 90 \end{bmatrix}
\end{aligned}
$$

Thus,

$$
A = \overbrace{\begin{bmatrix} 0.6667 & 0.1617 & -0.7276 \\ 0.6667 & -0.5659 & 0.4851 \\ 0.3333 & 0.8085 & 0.4851 \end{bmatrix}}^{Q} \underbrace{\begin{bmatrix} 3.0000 & 2.6667 & 2.6667 \\ 0 & 1.3744 & 0.6468 \\ 0 & 0 & 1.2127 \end{bmatrix}}_{R}
$$

As was mentioned earlier in 1.2 in veiw of change of basis and computing eigenvectors, in practice we prefer orthonormal matrices to avoid the expensive computation of matrix inverse, and use transposition instead.

Given $A = Q_1 R$ to begin with, we first compute $A_1 = RQ_1$. Since $R = Q_1^T A$, A_1 becomes *similar* to A. Consequently, we have the following iterations.

$$
\text{Let } A_1 = RQ \;\; = \;\; Q_1^T A Q_1
$$

$$\begin{aligned} \text{followed by } A_2 &= Q_2^T A_1 Q_2 \\ &\vdots \\ A_i &= Q_i^T A_{i-1} Q_i \\ &= (Q_1 Q_2 \cdots Q_i)^T A (Q_1 Q_2 \cdots Q_i) \end{aligned} \tag{67}$$

Example 11: Let

$$A = \begin{bmatrix} 0 & 1 \\ -2 & -3 \end{bmatrix}$$

We first apply the QR algorithm to obtain

$$Q_1 = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \quad \text{and} \quad R = \begin{bmatrix} 2 & 3 \\ 0 & 1 \end{bmatrix}$$

Consequently we get

$$A_1 = RQ_1 = \begin{bmatrix} -3 & 2 \\ -1 & 0 \end{bmatrix}$$

which is similar to A. If we repeat this procedure we get the following:

$$A_2 = \begin{bmatrix} -2.4000 & 2.8000 \\ -0.2000 & -0.6000 \end{bmatrix}, \quad A_3 = \begin{bmatrix} -2.1724 & 2.9310 \\ -0.0690 & -0.8276 \end{bmatrix},$$

$$A_4 = \begin{bmatrix} -2.0803 & 2.9708 \\ -0.0292 & -0.9197 \end{bmatrix}, \quad A_5 = \begin{bmatrix} -2.0388 & 2.9865 \\ -0.0135 & -0.9612 \end{bmatrix},$$

$$A_6 = \begin{bmatrix} -2.0191 & 2.9935 \\ -0.0065 & -0.9809 \end{bmatrix}, \quad A_7 = \begin{bmatrix} -2.0095 & 2.9968 \\ -0.0032 & -0.9905 \end{bmatrix},$$

$$A_8 = \begin{bmatrix} -2.0047 & 2.9984 \\ -0.0016 & -0.9953 \end{bmatrix}, \quad A_9 = \begin{bmatrix} -2.0023 & 2.9992 \\ -0.0008 & -0.9977 \end{bmatrix},$$

$$A_{10} = \begin{bmatrix} -2.0012 & 2.9996 \\ -0.0004 & -0.9988 \end{bmatrix}, \quad A_{11} = \begin{bmatrix} -2.0006 & 2.9998 \\ -0.0002 & -0.9994 \end{bmatrix},$$

$$A_{12} = \begin{bmatrix} -2.0003 & 2.9999 \\ -0.0001 & -0.9997 \end{bmatrix}, \quad A_{13} = \begin{bmatrix} -2.0000 & 3.0000 \\ 0.0000 & -1.0000 \end{bmatrix}$$

Thus we get A_{13}, an upper triangular matrix, that is similar to A. Notice that the eigenvalues of A are the diagonal entries of A_{13}.

It may be observed that the inverse of any Q_i is readily obtained by simply transposing it. Nevertheless, arriving at a diagonal matrix Λ is far away in most cases and we practically settle with some A_i that is upper-triangular which typically takes $\mathcal{O}\left(n^3\right)$ (i.e., not exceeding n^3) multiplications. There are several other issues, e.g., faster convergence, accuracy etc., which led to the development of variants of the algorithm. Interested reader may find those topics in the references.

The resulting upper triangular matrix A_i is similar to the original matrix A, and hence has the same eigenvalues, now sitting along the diagonal. Thus, without solving the characteristic polynomial we will be able to compute the eigenvalues spending only about n^3 saxpy floating point operations.

We will next see how this idea of decomposition can be extended to rectangular matrices.

1.7 Singular Value Decomposition

Let us now look at a matrix $A \in \Re^{m\times n}$, whose rank is $m\,(< n)$. There are plenty of applications where the linear transformations of this nature are necessary. Nevertheless, we proceed in the following manner building up creatively from our earlier experience.

1. We will first obtain the product $A^TA \in \Re^{n\times n}$, a symmetric matrix. If we perform a similarity transformation on this matrix, we first expect the diagonal matrix

$$\sum = \text{diag}\ \ (\sigma_1, \sigma_2, \cdots, \sigma_m,\ \sigma_{m+1} = 0, \cdots, \sigma_n = 0) \tag{68}$$

where we use the symbol σ for the eigenvalues of A^TA, which are better known as the *singular values* of A. Since the rank of A is $m < n$, the singular values $\sigma_{m+1}, \cdots, \sigma_n$ are zeros. It is standard practice to make this diagonal matrix with entries in the descending order, i.e., $\sigma_1 \geq \sigma_2 \geq \cdots \sigma_m > 0 = \sigma_{m+1} = \cdots = \sigma_n$.

For convenience, we consider the existence of an orthonormal matrix Q such that

$$Q^T\left(A^TA\right)Q\ = \sum = \ \text{diag}\ \ (\sigma_1, \sigma_2, \cdots, \sigma_n) \tag{69}$$

2. Left-multiplying on both sides by AQ, we get the following

$$(AQ)\cdot Q^TA^TAQ = AQ\sum$$

since Q is orthonormal we open up the paranthesis and initially obtain

$$\left(AA^T\right)AQ = AQ\sum$$

from which we readily obtain, by post-multiplying on both sides by Q^T

$$\begin{aligned} A &= \left(AA^T\right)^{-1}AQ\sum Q^T \\ &\triangleq U\sum V^T \end{aligned} \tag{70}$$

where V is the original Q itself, and U is obtained by pre-multiplying Q with $\left(AA^T\right)^{-1} A$. One should readily see that AA^T has full rank m, and hence it is invertible.

This is the basic idea of a singular value decomposition, popularly known as SVD.

If matrices are seen as generalizations of numbers, then we need the notion of the sign of a matrix and we briefly introduce the concept in the following section.

1.8 Sign Definite Matrices

Quite often, on a large scale, it becomes imperative to study quadratic polynomials of the form

$$\mathcal{P}(x) = \sum_{i,j=1}^{n} p_{ij} x_i x_j = p_{11} x_1^2 + \cdots + p_{nn} x_n^2 + p_{12} x_1 x_2 + \cdots + p_{n-1,n} x_{n-1} x_n \quad (71)$$

where $x = [x_1 \cdots x_n]^T$ is the vector of the variables x_i. Given x, $\mathcal{P}(x)$ evaluates to a scalar, typically a real number.

A general interpretion given to such polynomials is *energy* of a system; one can also see it, geometrically, as distance. Since energy or distance is a scalar that can never be negative irrespective of the sign of x_i, it is instructive to study the coefficients p_{ij}. In particular, it is easy to say that it is necessary to have $p_{ii} \geq 0$ as x_i^2 is always non-negative, but how about the other coefficients?

We may, with some effort, rewrite the quadratic polynomial as

$$\begin{aligned} \mathcal{P} &= x^T P x \\ &= x^T \begin{bmatrix} \sum_{j=1}^{n} p_{1j} x_j \\ \vdots \\ \sum_j p_{nj} x_j \end{bmatrix} \\ &= \sum_{i \leq j=1}^{n} p_{ij} x_i x_j \end{aligned} \quad (72)$$

where the matrix P may be observed to be a symmetric matrix; even if it were not symmetric, it does not matter as we can always get a symmetric matrix $S = \frac{1}{2}\left(P^T + P\right)$, which implies $x^T S x = x^T P x$. Thus no loss of generality is obtained by assuming P is symmetric.

Suppose we are able to get a vector z and an orthonormal matrix M such that $x = Mz$, and such that $M^T PM$ is diagonal then

$$\mathcal{P}(z) = z^T M^T PMz = \lambda_1 z_1^2 + \lambda_2 z_2^2 + \cdots + \lambda_n z_n^2 \tag{73}$$

containing only the square terms, which are obviously non-negative if and only if $\lambda_i \not< 0$. Ergo, the polynomial $\mathcal{P}(z)$ evaluates to a positive real number for any $z \neq 0$ if and only if all the eigenvalues λ_i of the symmetric matrix P are strictly positive. Extending this observation, we have the following five-fold classification of symmetric matrices

known as	eigenvalues
positive definite	$\lambda_i > 0$
positive semi-definite	$\lambda_i \geq 0$
negative definite	$\lambda_i < 0$
negative semi-definite	$\lambda_i \leq 0$
indefinite	some $\lambda_i \geq 0$ and others < 0

There is an interesting way, called Sylvester's criterion, which helps us identify the class of a given symmetric matrix. We will have a quick look at this criterion for a negative definite matrix. Since all the eigenvalues must be strictly negative,

$$-\lambda_1 > 0, \quad \lambda_1\lambda_2 > 0, \quad \ldots, \; (-1)^k \lambda_1\lambda_2 \cdots \lambda_k > 0, \quad \ldots$$

and each term may be identified as a part determinant: $\lambda_1\lambda_2 \cdots \lambda_k$ is the determinant of the matrix formed using the first k rows and k columns of P. One would readily notice that the eigenvalues are progressively covered in these sub-matrices, appropriately termed as *principal minors* earlier. Needless to say, for a positive definite matrix all the principal minors are strictly positive.

A couple of observations are in place here. First, like any norm, $x^T Px = 0 \Leftrightarrow x = 0$. However, $x^T Px = 0$ under certain conditions as well; for instance $(x_1 + x_2)^2 = 0$ if $x_2 = -x_1 (\neq 0)$. Thus, we often classify symmetric matrices as either positive semi-definite, or negative definite, or indefinite. Secondly, looking at the variables x_i as being statistically available with probabilities p_i, we obtain the covariance matrices as symmetric positive definite matrices which might be interpreted as squares of *weighted* 2-norms, i.e.,

$$x^T Px = \|x\|_P^2 \tag{74}$$

We have one last piece to explore before we close this appendix.

1.9 The Condition Number

The argument we present here, nothing to do with the definitions, is that we are *perturbing* the diagonal elements to check how far away the transformation is from being a bijection. The range of eigenvalues could be small or large; in the later case, we end up with ill-conditioned systems and such concepts, which are more meaningful from this perspective - there are eigenvalues close to zero which create troubles in the computation. The condition number of a matrix is given by

$$\kappa(A) \;=\; \|A\| \cdot \|A^{-1}\| \tag{75}$$

which indicates that $\kappa(A) > 1$, but if it is too much bijection of the transformation could be jeopardized. If we imagine a diagonal matrix with eigenvalues in the interval $[\lambda_{max}, \lambda_{min}]$,

$$\|\Lambda\| = |\lambda_{max}|, \quad \|\Lambda^{-1}\| = \left|\frac{1}{\lambda_{min}}\right| \quad \text{and} \quad \kappa(A) = \left|\frac{\lambda_{max}}{\lambda_{min}}\right| \tag{76}$$

The implication is that, if we have a wide range of eigenvalues, with one close to zero and one reasonably large in magnitude, the condition number is going to be quite large. In the solution to the $Ax = b$ the eigenvalue relatively close to zero plays spoil sport as finite precision arithmetic is likely to round it off to zero. We illustrate this in the following example.

Example 12: Let us solve the linear system

$$A = \begin{bmatrix} 1 & 1 \\ 0.49 & 0.51 \end{bmatrix}, \quad b = \begin{bmatrix} 2 \\ 1 \end{bmatrix}$$

Without any calculation one would get the solution as $x = [1 \quad 1]^T$ with $\|x\|_2 = \sqrt{2}$. But, if we slightly change b by adding an ϵ to 2,

$$x \;=\; \begin{bmatrix} 1 + 25.5\epsilon \\ 1 - 24.5\epsilon \end{bmatrix}, \qquad \|x\|_2 = \sqrt{2 + 2\epsilon + 1250.5\epsilon^2}$$

The reason is that the eigenvalues of A are

$$\lambda_{1,2} = 1.496636703514598 \approx 1.50, \quad \text{and} \quad 0.013363296485402 \approx 0.01$$

significantly far apart from each other, and $\kappa(A) = 125.0020$.

What is Ahead?

With these pages of development, the next question is: "where do we go now?" For a wholesome development a sequence of right questions is important, with hopefully a big picture sketched. This was the chief intention of this note. Once we identify a vector and its space $\mathcal{V}^n$ over an arbitrary field $\mathcal{F}$, for instance, a continuous real-valued function f on a closed interval I, we quickly look at αf and hence $\alpha f + \beta g$. If we look at objects this way, we find that there is more structure showcasing extra features which distinguish it from others. The strength of this kind of an abstract approach lies in the fact that consequences of the axioms can be applied to many different examples; one may instantaneously recall that differentiation (or integration), Laplace transformation, and even computing the expectation of a random variable follow the suit.

This is the beginning of a formal course on linear algebra and we direct the reader to the following excellent references.

1.10 References

1. J.B. Fraleigh and R.A. Beauregard, *Linear Algebra*, 3/e, Addison-Wesley, 1995.

2. C.D. Meyer, *Matrix Analysis and Applied Linear Algebra*, 2/e, SIAM, 2023.

3. D. Poole, *Linear Algebra: A Modern Introduction*, 4^{th} ed., Cengage Learning India, 2014.

4. M. Artin, *Algebra*, Pearson Education, 2/e (updated), 2024.

5. P.M. Cohn, *Classic Algebra*, John Wiley & Sons, 2000.

6. G.H. Golub & C.F. Van Loan, *Matrix Computations*, 4/e, Johns Hopkins Univ. Press, 2013.

7. Nicholas Higham, *Functions of Matrices: Theory & Computation*, SIAM, 2008.

8. Nicholas Higham, *Accuracy and Stability of Numerical Algorithms*, 2/e, SIAM, 2002.

9. Paul R. Halmos, *Finite Dimensional Vector Spaces*, 2/e, Van Nostrand, 1958