Guide to Part I.

A *matrix* is simply a rectangular array of numbers, in itself nothing mysterious. However, the algebraic interactions of such arrays, and their geometric manifestations, open up a world of fascinating mathematical phenomena. This is matrix theory, also known as linear algebra.

To keep things simple, Part I of this course will limit itself by and large to just two kinds of matrices:

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \qquad \text{and} \quad X = \begin{bmatrix} x \\ y \end{bmatrix}.$$

When we say "matrix", we shall usually mean the first of these types, preferring the term "vector" for the second. To save space, we shall often write (x, y) instead of the vertical column displayed above when dealing with vectors.

Vector quantities abound. For instance, instead of considering an investor's assets as a single amount, you might break it down into separate values for stocks and for bonds; or, instead of looking at a total population of (say) eagles, you might count adults and sexually immature juveniles as separate subpopulations. To study how these numbers change from one year to the next, you generally have to consider four coefficients: the growth rates of x and y by themselves, as well as the factors by which they hinder or enhance each other. This is where matrices come in; most frequently they occur as *multipliers of vector quantities.* The multiplication of the vector X by the matrix A shown above is defined as follows:

$$AX = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} ax + by \\ cx + dy \end{bmatrix}.$$

If the terms b and c are both $= 0$, the matrix A is called *diagonal*. In that case, the components of AX are just ax and dy with no cross-terms, and the problem does not really need vectors; it is just one ordinary 1-dimensional problem on top of another. One of the major themes of this course is the search for ways of simulating this happy state of affairs ("diagonalization").

1. Thus far, the so-called multiplication by A defined above must seem like a pompous way of rewriting a couple of harmless linear formulas $ax + by$ and $cx + dy$. What really makes it fly is its application to a *pair* of columns, i.e. to another 2×2-matrix B. The resulting product is again a 2×2-matrix $C = AB$. So, square matrices begin to look like a kind of super number-system. Indeed, the usual associative and distributive laws hold, associativity is less obvious than one might hope. In fact, *it* is what motivates the whole row-column alternance so typical of matrix multiplication. Commutativity fails utterly, witness

$$\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \neq \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}. \tag{1}$$

The fourth from the left is called the *zero matrix*, for obvious reasons.

We are interested in these rules of algebra, because we are working toward a symbolic language, in which matrices appear as single entities, not as collections of numbers. To handle formulas in this language, we must know how these entities may or may not be shuffled around.

2. Since we can multiply by the matrix A, it would be comforting to know that we can also "divide" by it. Is there a matrix A^{-1} which undoes multiplication by A, just as $1/3$ cancels the multiplication by 3? The answer is *yes* for the first and *no* for the second of the following matrices

$$\begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix} \qquad \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}. \tag{2}$$

In general it is yes unless A is *singular*, which means that it can behave like a zero multiplier (AB can be zero, even if B is not), but does *not* mean (as with ordinary numbers) that A is necessarily zero.

This lesson describes several tests for singularity and a method for finding A^{-1} when it exists. Much of it involves the number $(ad - bc)$, called the *determinant* of A.

Since left and right multiplication by A are not always the same, there is no immediate reason for left and right inverses to be equal. They do coincide in theory, but not always in computational practice.

3. Geometrically a single real number x does not have much character except for its size. However, a *pair* (x, y) of numbers shows up as a point or a vector in the plane. Whether it should be pictured as one or the other is a secondary question: the numbers themselves do not care how we view them.

A square matrix A can be visualized as transforming these points or vectors by multiplication. Using three specific types of matrices, we try to develop some feeling for such deformations of the plane. These three types are exemplified by the matrices

$$\begin{bmatrix} 4/5 & 0 \\ 0 & 3/2 \end{bmatrix}, \qquad \begin{bmatrix} 1 & 1/2 \\ 0 & 1 \end{bmatrix}, \qquad \begin{bmatrix} \sqrt{3}/2 & -1/2 \\ 1/2 & \sqrt{3}/2 \end{bmatrix}, \tag{3}$$

which transform the unit square (dotted lines) in the manner indicated by the following sketches.

These three archetypes (diagonal, shear, and rotation) will be met again and again. We shall eventually see that *any* matrix can be expressed as a product of shears and diagonals, or of rotations and diagonals. In the more immediate future, every 2×2-matrix will be made to appear "similar" to a scalar multiple of just one of these types.

4. Matrices are more difficult to size up than single numbers, since they may look more or less alike and yet be very different in nature, as for instance,

$$\begin{bmatrix} 6 & -2 \\ 2 & 1 \end{bmatrix}, \qquad \begin{bmatrix} 5 & -2 \\ 2 & 1 \end{bmatrix}, \qquad \begin{bmatrix} 5 & -3 \\ 3 & 1 \end{bmatrix}. \tag{4}$$

The best way of analysing their true character is to look for "eigenlines", that is, lines through the origin which stay put as the plane gets transformed by A. Within such a line Λ, transformation by A has the effect of multiplying by a scalar λ, the so-called *eigenvalue* corresponding to Λ.

The eigenvalues of A (if any) must satisfy a certain quadratic equation and thereby are easily tracked down. Once they are known, it is child's play to determine the corresponding eigenlines. Of the three matrices shown in (4), the first has two eigenvalues (5 and 2) and two eigenlines (generated by the vectors $(2, 1)$ and $(1, 2)$, respectively). The second has the single eigenvalue 3 with eigenline $x = y$; the third has no eigenlines.

The same three patterns are exhibited by the archetypal matrices of the last paragraph: the diagonal has two eigenlines (the coordinate axes), the shear has just one (with eigenvalue 1), and the rotation has none. A matrix with two eigenlines is called *diagonalizable.* Such matrices are literally *similar* to diagonal ones: if you take the eigenlines as axes of a skewed coordinate system, the resulting new coordinates really experience the matrix action as being diagonal. An explanation in terms of matrices will be given below. How this might be manifested in reality is described in the "money story" at the end of this guide.

5. The technical definition of similarity between the matrices A and B is that $B = M^{-1}AM$ or equivalently $A = L^{-1}BL$, for suitable invertible matrices M and $L = M^{-1}$. If you want to look at similarity as being a kind of change of coordinates, as hinted above, these matrices act as "dictionaries" for translating from the original coordinates to the new ones (by L) or vice versa (by M). Thus, every X has the "alias" LX, the "new" coordinates of X. The similarity relation

$$L(AX) = B(LX) \tag{5a}$$

explicitly says: the alias of AX is obtained by applying B to the alias of X. Thus, in the world of aliases, B takes the place of A.

We are particularly interested in the case where B can be made diagonal; i.e. where A has two eigenlines. But even if this is not so, it is still possible to relate A to one of the three archetypes. For the matrices shown in (4), these "standard forms" are

$$\begin{bmatrix} 5 & 0 \\ 0 & 2 \end{bmatrix}, \qquad \begin{bmatrix} 3 & 1 \\ 0 & 3 \end{bmatrix}, \qquad \begin{bmatrix} 3 & -\sqrt{5} \\ \sqrt{5} & 3 \end{bmatrix}. \tag{5b}$$

The second of these is a shear multiplied by 3, the third is a rotation multiplied by $\sqrt{14}$. The translation matrix M is always computable from information coming out of the search for eigenlines. If A is diagonalizable, the columns of M are just eigenvectors.

6. The diagrams at the end of paragraph 3 are trying to suggest how A deforms the entire plane by a single multiplication applied to all points at once. Now we are preparing for another kind of picture, one which shows the movement of a few typical points as they are being hit by A over and over again. Given an initial point X, we want to investigate the behaviour of its "orbit" $A^n X$ as $n \to \infty$. This is analogous to but more interesting than the sequence $a^n x$ of successive population sizes in ordinary exponential growth. The "bird story" at the end of this guide describes how it might relate to reality.

The lesson starts by explaining why we can shift the investigation to the standard forms B (types (a) to (c)) instead of staying with the original A. If B happens to be a rotation, $B^n X$ will of course go round and round, literally staying in orbit. But this is rare; for the general B of type (c), the sequence $B^n X$ will lie on a spiral. In cases (a) and (b) the trend is for the typical "orbit" to be attracted to an eigenline. We can see a hint of this behaviour by watching the diagonal of the unit square in the pictures of paragraph 3. For numerical evidence we turn to to the matrices (4), whose eighth powers are as follows:

$$\begin{bmatrix} 5.2 & -2.6 \\ 2.6 & -1.3 \end{bmatrix}, \qquad \begin{bmatrix} 4.2 & -3.5 \\ 3.5 & -2.8 \end{bmatrix}, \qquad \begin{bmatrix} -.16 & .47 \\ -.47 & .47 \end{bmatrix}, \tag{6}$$

each multiplied by 10^5. Both columns of the first matrix are so close to the eigenline $x = 2y$ that they already appear to lie on it; this puts $A^n X$ (approximately) on that line for *any* X. The second matrix, which is related to a shear, approaches the expected eigenline $x = y$ much more slowly. The third matrix produces outward spirals and does not settle on any particular direction.

7. This lesson surveys and classifies the shapes of orbits produced by the different kinds of matrices. The following portraits pertain to matrices with one or no eigenlines (cases (b) and (c)). The diagonalizable case is shown in the text itself.

For the sake of clarity we have represented the orbits as continuous curves rather than sequences of dots. We exclude matrices with negative eigenvalues, which would lead to confusing zig-zags; in practice they can always be avoided by taking only *even* powers. Other phenomena which remain undepicted are the ones caused by eigenvalues 0 (collapse to a line) and 1 (stationary points). They would make lousy pictures.

8. So far the only direct interaction between two vectors has been via addition. Now we thicken the plot by a kind of "multiplication" for vectors, in which the "product" is not another vector but a scalar. This new ingredient is so basic that we could have introduced it at the very beginning of the course. Accordingly, this lesson does not depend on any of the previous ones.

The formal algebraic properties of the "dot product" are:

$$V \bullet X = X \bullet V, \qquad V \bullet (X + \alpha Y) = V \bullet X + \alpha(V \bullet Y), \qquad \text{and} \quad V \bullet V > 0, \tag{8a}$$

for $V > 0$. Using just these, we define the *norm* of a vector, the *orthogonality* of a vector-pair, and the *projection* of one vector onto another. The latter has an interesting minimum property.

Now comes the surprise: by a 2500 year old theorem, the formmula $|V|^2 = a^2 + b^2$ shows that $|V|$ is the *length* of the arrow representing V in rectangular coordinates. It quickly follows that the dot-product $V \bullet X$, which at first looked like just another algebraic gimmick, has a very concrete geometric meaning, namely $|X||V| \cos \phi$, where ϕ stands for the angle between the two vectors.

To illustrate the power of this, consider the triangle formed by the points $P = (-3, 2)$, $Q = (1, 5)$, and $R = (9, 7)$. If you want to find the angle QPR, you first note that the vectors $V = (4, 3)$ and $W = (12, 5)$ take you from P to Q and from P to R, respectively. Then

$$|V|^2 = 4^2 + 3^2, \qquad |W|^2 = 12^2 + 5^2, \qquad \text{and} \quad V \bullet W = 4 \cdot 12 + 3 \cdot 5 \tag{8b}$$

show that $\cos \phi = 63/65$. Looking up the inverse cosine, you find $\phi = 14^\circ\ 15'$.

9. As soon as the dot product gets involved, the flavour of matrix theory becomes decidedly more geometric. Two kind of matrices are particularly well-behaved toward the dot product: the *orthogonal* and the *symmetric* ones. Their good behaviour is characterized by the equations

$$AV \bullet AW = V \bullet W \qquad \text{and} \qquad V \bullet AW = AV \bullet W, \tag{9a}$$

respectively. Both of these qualities can also be expressed in terms of the transpose A^T, and both of them have geometric meanings.

Orthogonal matrices are either rotations through some angle or reflections across some line. Symmetric matrices are always diagonalizable, their two eigenlines are perpendicular to each other. Such a matrix S can therefore be diagonalized by "conjugation" with a suitable *orthogonal* Q, i.e. $S = QDQ^{-1}$.

It turns out that *any* matrix A can be written as $A = SR$, with S symmetric and R orthogonal. This is known as the *polar decomposition* of A. If we write $S = QDQ^{-1}$ and let $P = Q^{-1}R$, this assumes the form $A = QDP$, which is called the *singular value decomposition* and shows A as the product of a diagonal sandwiched between two orthogonals. By changing a sign in D (if necessary), P and Q can always be chosen to be rotations. It may be hard to imagine this for a shear. The equation

$$\frac{1}{\sqrt{1+k^2}} \begin{bmatrix} 1+2k^2 & k \\ k & 1 \end{bmatrix} \cdot \frac{1}{\sqrt{1+k^2}} \begin{bmatrix} 1 & k \\ -k & 1 \end{bmatrix} = \begin{bmatrix} 1 & 2k \\ 0 & 1 \end{bmatrix} \tag{9b}$$

shows how it is done.

10. The last lesson of Part I could easily have been the first or second one. It deals with an alternative approach to inverses and determinants, called "elimination" or "row reduction" and often dignified with one of the greatest names in mathematics — Gauss. The basic idea is familiar to anyone who has ever solved linear equations in several variables: you add and subtract the equations to and from each other with the aim of eliminating as many variables as possible.

Let us watch this process being applied to a matrix A with entries 1,2, 3, and 4.

$$A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} \xrightarrow{E_1} \begin{bmatrix} 1 & 2 \\ 0 & -2 \end{bmatrix} \xrightarrow{E_2} \begin{bmatrix} 1 & 2 \\ 0 & 1 \end{bmatrix} \xrightarrow{E_3} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = I. \tag{10a}$$

Naively, the steps E_1 to E_3 are "elementary row operations": E_1 subtracts 3 times row one from row two, E_2 multiplies row two by $-1/2$, and E_3 subtracts 2 times row two from row one. From a more sophisticated point of view, each operation amounts to left multiplication by a certain "elementary matrix". E_1 and E_3 represent shears, and E_2 a diagonal matrix. Thus (10a) can be rewritten as

$$G_{12}(-2)D_2(-1/2)G_{21}(-3)A = I \qquad \text{or} \qquad A = G_{21}(3)D_2(-2)G_{12}(2). \tag{10b}$$

All sorts of questions involving A can be answered from this factorization. In particular, the determinant is concentrated in the D-type factors; in this case it is -2.

As a counterpart to the polar decomposition for a shear, you might be interested to see how a rotation decomposes into shears and diagonals: $R(\theta) = G_{12}(s/c)D(c, 1/c)G_{12}(-s/c)$, where $s = \sin\theta$ and $c = \cos\theta \neq 0$.

****

Synopsis. To get an overview of this theory it may be helpful to think of it as having three facets often overlapping in several ways:

* *Elimination - Inversion*, lessons 2 and 10 with preliminaries in lessons 1 and 3.
* *Eigenvectors - Similarity*, lessons 4 and 5 with applications in lessons 6 and 7.
* *Dot Product - Orthogonality*, lessons 8 and 9.

Of the three items listed above, the middle one is the most difficult. It has been given the largest spread here in Part I, because it is already interesting and still easily manageable in two dimensions. The more geometric last item reaches the right balance of challenge versus tractability in three dimensions, and will therefore be stressed in Part II. Our main focus in Part III will be the first item, which is conceptually the simplest of the three.

Operationally, we usually manipulate a given matrix A in one of two ways, depending on the problem to be solved: (left) *multiplication* to make a simpler $A' = MA$, or *conjugation* to get a friendlier $B = M^{-1}AM$. The first of these is also related to the "factorization" or "decomposition" $A = M^{-1}A'$ in which both factors are presumably more workable than the original A. The invertible matrix M is sometimes built up by easy instalments (as in elimination), and at other times carefully constructed in one piece (as in the polar decomposition).

Story # 1: Magic Monies.

This is the story of a ficticious Cuban capitalist, a lady named Ahorita Vengo, who holds considerable sums of both Canadian and American dollars. She avoids money conversion muddles by doing all her calculations in Cuban pesos instead of dollars. If x and y denote the number of millions of pesos invested in Canadian and American currencies, respectively, the annual growth of these monies can be represented by the matrix multiplication

$$X \quad \mapsto \quad AX = \begin{bmatrix} 1.09 & -.04 \\ \\ -.03 & 1.05 \end{bmatrix} \begin{bmatrix} x \\ \\ y \end{bmatrix} = \begin{bmatrix} 1.09x - .04y \\ \\ -.03x + 1.05y \end{bmatrix}. \tag{1}$$

Let us not worry right now about the derivation of this equation. (For the curious, here is how it works: firstly, there is a higher annual interest rate in Canada (12%) than in the US (9%); secondly, the Cuban government taxes Canadian holdings at a lesser yearly rate (6%) than the American ones (8%); thirdly, these taxes on North American accounts are payable half in Canadian dollars and half in US dollars, irrespective of their source.)

Ahorita is irritated by the irregularity of this growth pattern. If she puts 3 million pesos into Canadian ($x = 3$) and an equal amount into American ($y = 3$), she winds up, after one year, with 3.15 in one currency and 3.06 in the the other. Unfortunately this does not mean that x consistently grows at 5% and y at 2%. If she had split her 6 million by putting $x = 4$ and $y = 2$, the results would have been 4.26 and 1.98, respectively, suggesting growth rates of 6.5% and −1%. Not even the totals are the same: the 6 would have grown to 6.21 under the first scheme and to 6.24 under the second.

Her sister-in-law, Megusta Poco, is a professional broker and tells her about the magic $40 - 60$ split, suggesting that she invest 2 million in Canadian and 3 million in US money. Sure enough, if she puts $x = 2$ and $y = 3$, she gets 2.06 and 3.09 for the next year, a clean 3% gain across the board. Best of all, the $2 : 3$ ratio is maintained, and the same percentage increase will be repeated every year, great! She leaves the 5 million in Poco's care and turns her attention to other matters.

She still has 1 million left to play with, so a few days later she goes to consult Poco's arch-rival Loquiero Mucho, a somewhat disreputable old dude, feared and respected for his uncanny knack of making money. Mucho says that there is another magic ratio, and on receipt of a hefty consultation fee divulges it as $2 : -1$. What does this mean? Take your million, Mucho says, borrow another million in US funds (this makes $y = -1$), and put the 2 million into Canadian dollars (so $x = 2$). Lo and behold: Ahorita gets 2.22 for x and −1.11 for y after one year, the same 11% in both components and, of course, for the total. Again the ratio is maintained, and the same fabulous return will swell her account year after year.

Ms. Vengo's first impulse is to take all her money out of Poco's hands and give it to Mucho, but then she remembers her family ties and also the nasty law which forbids her making a net debt in a foreign currency. At least she is pleased with the neatness of the new arrangement. Instead of getting confused in the vagaries of Canadian versus American dollars, she now thinks in terms of Mucho and Poco portfolios. The former consistently grow at 11%, the latter at 3%. It is as if she were working in a new monetary coordinate system — and with a diagonal matrix.

In fact, Vengo's brokers have stumbled onto the two ratios $2 : -1$ and $2 : 3$ defining the *eigenlines* of the matrix A, which allow a "diagonal" reorganization of her problem.

Story # 2: Birds and Bunnies.

In many species, population growth can be studied by keeping track of two types of females: *adults* and *juveniles*. Juveniles are those who were born in the last breeding season and survived to the present one. Presumably they produce less offspring than adults (often none at all) and survive less frequently to the next breeding season (in which they will be adults). Males are ignored in this kind of analysis, apart from being assumed numerous enough to fill out the reproductive capacities of the females.

Let x_n be the the number of adults, y_n the number of juveniles, and $t_n = x_n + y_n$ the total number of females (excluding babies) in the n-th season. What will be the situation in the following season? If a and b denote the survival rates for adults and juveniles, respectively, the number of adults will be $x_{n+1} = ax_n + by_n$. If every adult contributes (on the average) c juveniles to the next season, and every juvenile likewise contributes d future juveniles (usually $d = 0$), the younger set will number $y_{n+1} = cx_n + dy_n$. This can be summarized in the matrix equation

$$\begin{bmatrix} x_{n+1} \\ \\ y_{n+1} \end{bmatrix} = \begin{bmatrix} a & b \\ \\ c & d \end{bmatrix} \begin{bmatrix} x_n \\ \\ y_n \end{bmatrix}.$$

Abbreviating the square matrix by A and the n-th "population vector" by X_n, this simplifies to $X_{n+1} = AX_n$, hence $X_n = A^n X_0$.

For any A with $a, b, c > 0$, and $d \geq 0$, the evolution of a non-negative population vector $X_n = A^n X_0$ is determined by two key facts:

1. As n increases, the fraction x_n/t_n approaches a constant x^* independent of the initial x_0, y_0. Of course, the y-fraction y_n/t_n will simultaneously approach $y^* = 1 - x^*$. We call this the *stable distribution*. Any X with $x/t = x^*$ and $y/t = y^*$ will be called a *stable vector*.
2. If X is a stable vector then so is AX. This means that, on a stable vector, multiplication by A has a very simple effect. Since the ratio between the components cannot change, both of them get multiplied by the same *growth factor* $\lambda > 0$. So, a stable X "experiences" A as if it were a simple number instead of a matrix: $AX = \lambda X$.

Taken together, these facts tell us that, if we wait long enough, X_n will approach stability and stay there. From then on our problem reduces to the familiar one-dimensional growth (or decay) of $t_n = x_n + y_n$ by the factor $\lambda > 0$. Our efforts must therefore be directed toward finding the stable distribution and the growth factor.

The most obvious (though somewhat tedious) method is to look at increasing powers of A until we see the same ratio jelling in both columns. To get high powers of A quickly, it is advisable to square over and over, forming $A^2 = AA, A^4 = (A^2)^2, A^8 = (A^4)^2, A^{16} = (A^8)^2$, and so on.

In lesson 4 below you will learn a short-cut which directly computes the growth factor (also called "eigenvalue") as

$$\lambda = \frac{1}{2}\left[(a + d) + \sqrt{(a - d)^2 + 4bc}\right].$$

Once found, this number is easily used to solve the linear equation $(A - \lambda I)X^* = 0$ subject to the condition that $x^* + y^* = 1$. Actually the naive method of forming powers of A is not as foolish as it might seem: many practical algorithms for finding eigenvalues of larger matrices are based on modifications of the power method.

Example 1: For some species of birds, the reproduction story outlined above might yield a matrix with $a = .7,\ b = .3,\ c = 2,\ d = 0$. Repeated squaring yields

$$A^2 = \begin{bmatrix} 1.09 & 0.21 \\ 1.40 & 0.60 \end{bmatrix} \quad A^4 = \begin{bmatrix} 1.48 & 0.35 \\ 2.37 & 0.65 \end{bmatrix} \quad A^8 = \begin{bmatrix} 3.04 & 0.76 \\ 5.05 & 1.27 \end{bmatrix}.$$

From the last power we see that the x-part of both columns seems to go toward $3/8 = .375$. This is borne out by checking A^{16}, the square of A^8. Thus $x = 3, y = 5$ should give a stable vector. Indeed multiplication by A yields $(.7)(3) + (.3)(5) = 3.6$ and $(2)(3) + (0)(5) = 6$, which is in the same ratio $3 : 5$ but multiplied by $\lambda = 1.2$

Example 2: In a rare species of gnomes, the mating game is played year after year, as follows: .8 of married gnomes stay married for another year, .2 split up; .7 of single gnomes remain alone for another year, .3 get married. The matrix which governs this madness has $a = .8,\ b = .3,\ c = .2,\ d = .7$. Proceeding as before, we get

$$A^2 = \begin{bmatrix} .7 & .45 \\ .3 & .55 \end{bmatrix} \quad A^4 = \begin{bmatrix} .625 & .562 \\ .375 & .438 \end{bmatrix} \quad A^8 = \begin{bmatrix} .6016 & .598 \\ .3984 & .402 \end{bmatrix}.$$

It looks as though $x^* = .6,\ y^* = .4$ is the stable distribution. To test it, we multiply by A—- and get the same values back. Hence $\lambda = 1$. This one is a *stochastic* matrix, meaning that both columns sum to 1. Such matrices are an important but not very typical special case: for large n, the columns of A^n actually tend toward becoming *equal*, and λ is always 1, i.e. there is no over-all growth or decay.

Example 3: One of the earliest " modern" European mathematicians (around 1200 A.D.) was Leonardo da Pisa, also known as Fibonacci, the man who introduced Arabic numerals to the west. Today he is remembered mainly for the numerical sequence

$$1, 1, 2, 3, 5, 8, 13, 21, \ldots,$$

generated by the powers of the matrix A which has $a = b = c = 1$, and $d = 0$. These numbers have such weird and wonderful properties, that they are still the object of serious study. It is remarkable that Leonardo actually posed his original problem in the context of the reproduction of rabbits (immortal ones, since their survival rate is 1). Here we have

$$A^4 = \begin{bmatrix} 5 & 3 \\ 3 & 2 \end{bmatrix} \quad A^8 = \begin{bmatrix} 34 & 21 \\ 21 & 13 \end{bmatrix} \quad A^{16} = \begin{bmatrix} 1597 & 987 \\ 987 & 610 \end{bmatrix}.$$

In the 16-th power the ratios for the first and second columns are both equal to 1.61803. This number also happens to be the growth factor, whose theoretical value is $\lambda = (1 + \sqrt{5})/2$.

MATH 221: A Twelve Week Course.

In the following outline, each item corresponds to one week (3 hours) with one home-work assignment. The material has been arranged in such a way that the first two months are essentially restricted to two and three dimensions respectively. This has the advantage of concreteness and immediacy. The student will learn basic matrix lore and notation in a context where examples are compact and varied. The step up from two to three dimensions naturally motivates some of the subsequent techniques and considerations. At the same time the visual connection remains available for another few weeks.

The only addition to the traditional material is the application (week 3) to linear difference equations. The course *needs* an applied topic, and this one is natural. Time has been made for it by deflating the topic of week 12 somewhat, by treating linear equations less broadly, and by omitting an extensive discussion of lines and planes. (There is an obvious difference between including such matters in examples and exercises, and putting them on the official lecture menu.)

The numbers in brackets refer to the lessons in these pages.

1. The algebra of 2×2-matrices. Product, inverse, determinant; singularity vs. invertibility. Geometric view of matrix action, typical examples. *[1,2,3]*

2. Eigenvectors and similarity (2×2). Characteristic polynomial, diagonalizable matrices; three canonical forms. *[4,5]*

3. Powers and difference equations. Transition matrices in population models; linear differential equations via Euler's method. *[6,7]*

4. Dot product and matrix geometry. Dot in two and three dimensions. Symmetric and orthogonal matrices in the plane. *[8,9]*

---

5. Elimination and elementary matrices. Prologue in the plane. Invertibility in three dimensions; LU-factorization. *[10,11]*

6. Determinants (3×3) and cross products. Laplace expansion, connection with elementaries, multiplicativity. Geometry of determinant and cross product. *[12,17]*

7. Diagonalizability. Characteristic polynomial, multiplicity of eigenvalues; eigenlines and eigenplanes. *[13,14]*

8. Symmetric and orthogonal matrices. Spectral theorem (dim ≤ 3), rotations and reflections in 3-space. *[15,16]*

---

9. Independence and dimension. The fundamental theorem of square matrices; invariance of dimension. *[21,23]*

10. Row space and column space. Finding bases of various subspaces of $\mathbf{R}^n$ by elimination. Rank versus nullity. *[22,24]*

11. Orthogonalization. The Gram-Schmidt process; least squares; linear regression and curve fitting. *[18,25]*

12. Vector spaces and linear transformations. Introduction to the abstract picture; function spaces. *[ad lib.]*