Linear Algebra: A Happy Chance to Apply Mathematics

Gilbert Strang
Department of Mathematics, MIT
Cambridge MA 02139 USA
gs@math.mit.edu
math.mit.edu/∼ gs

I believe that linear algebra is *the* most important subject in college mathematics. Isaac Newton would not agree! But he isn't teaching mathematics in the 21st century (and maybe he wasn't a great teacher, we will give him the benefit of the doubt). Certainly Newton demonstrated that the laws of physics are best expressed by differential equations. He needed calculus: quite right. But the scope of science and engineering and management (and life) is now so much wider, and linear algebra has moved into a central place.

My presentation at ICME-10 had four parts. The first transparency was an extremely short outline, and I will follow it here:

1. **Spirit**
2. **Content**
3. **Web**
4. **Examples**

1. Spirit In my experience, the "spirit" of the class is all-important. Students need to know that the course is *prepared for them.* The best part of linear algebra is its wonderful mixture of ideas and applications, theory and practice, elegance and purpose. We are missing that best part if we present linear algebra as a dry and abstract subject—an axiomatic exercise. The beauty of this course is that the central ideas are really needed and used.

Good teaching is adapted to the students. You can have the pleasure of seeing a change in their (wrong) ideas about mathematics. Recognizing patterns is a natural human activity, for us to encourage in every student. Fortunately, the patterns of linear algebra (reflected in important matrices) are *not too hard to learn.*

2. Content May I briefly outline a typical course? The details will depend on the number of lectures (and also the starting point—students may have seen 2 by 2 and 3 by 3 matrices already). Even for small matrices, there is thinking to be done—here is a challenge exercise borrowed from my textbook [1]:

> If (a, b) is a multiple of (c, d), *show that* (a, c) *is a multiple of* (b, d). You could use numbers first to see how a, b, c, d are related, assuming $abcd \neq 0$. This question will lead to:
>
> If $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$ has dependent rows then it has dependent columns.
>
> If $AB = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ then $BA = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$. That looks so simple. . .

This exercise illustrates how ideas (and words) can be introduced early. When the precise definitions arrive, the intuition is there to meet them. We all understand by examples!

The crucial concepts of linear independence and dependence and rank are "foreshadowed" by working with a 2 by 2 matrix. That size is also just right for eigenvalues, to connect to the trace (λ_1 *plus* λ_2) and also the determinant (λ_1 *times* λ_2). Before this must come $Ax = b$ and the elimination algorithm that simplifies all linear systems. Here is a course outline.

$Ax = b$ by elimination, leading to triangular factors $A = LU$
Nullspace and column space of A and A^{T}
Linear independence, basis, and dimension
    Orthogonality of these four fundamental subspaces
Least squares and $A^{\mathrm{T}}A\widehat{x} = A^{\mathrm{T}}b$
    Projections and orthonormal basis by Gram-Schmidt ($A = QR$)
Determinants: Properties and formulas
    Cramer's Rule, formula for A^{-1} and volume
Eigenvalues and Eigenvectors
    Powers A^k and $\frac{du}{dt} = Au$ by diagonalizing $A = S\Lambda S^{-1}$
    Symmetric matrices with $A = Q\Lambda Q^{\mathrm{T}}$
    Positive definite $A^{\mathrm{T}}A$ with pivots > 0 and eigenvalues > 0
    Singular value decomposition (SVD) $A = U\Sigma V^{\mathrm{T}}$
Linear transformations and change of basis
Linear algebra in engineering, economics, networks, ...

3. Web The course page web.mit.edu/18.06 has become a valuable link to the class, and a resource for the students:

1. Assignments and exam solutions and schedule changes are more secure on the web than in a student's mind (not to speak of the professor's mind).

2. The web page contains short applets with animated graphics, and also complete videos of the linear algebra lectures in an earlier year.

The video lectures use RealPlayer. They can be downloaded from ocw.mit.edu. I must say more about the animated graphics:

> The presentation of eigenvalues and eigenvectors now includes *voiceover*. I explain the graphics in words, while the vectors are moving. There are also 1-minute "mini-lectures"—the first goes from $Ax = \lambda x$ to $\det(A-\lambda I) = 0$.

I am very optimistic about the potential for graphics with sound. The bandwidth is low, and FlashPlayer is freely available. This offers a *quick review* (and active experiment) while the full lectures offer a more complete review.

The lectures (soon the voiceovers) are also on MIT's OpenCourseWare site ocw.mit.edu. The quality of the reception is improved by worldwide relays and by offering three different transmission rates. OCW has given strong support. The applied mathematics websites on ocw.mit.edu and math.mit.edu/18085 also have videos that begin with applied linear algebra. (The first lecture discusses the four matrices K, T, B, C from our examples below.)

The text for the MIT course 18.06 is my book *Introduction to Linear Algebra* [1]. Many teachers contact me about this book—I am very happy about its reception. Others know the earlier book *Linear Algebra and Its Applications.* Its first edition in 1976 helped to move linear algebra courses toward the audience that needed this subject. So many students respond to examples and applications, rather than axioms and proofs. This is good. The mathematics is still there. I hope the new 2005 edition [2] will open this subject to an even wider audience, to show the purpose of mathematics and its value *directly to students.*

4. Examples I will focus first on A *times* x, a matrix times a vector. It is essential to see Ax as a ***combination of the columns of*** A:

$$Ax = \begin{bmatrix} 1 & 3 \\ 2 & 5 \\ 4 & 9 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = x_1 \begin{bmatrix} 1 \\ 2 \\ 4 \end{bmatrix} + x_2 \begin{bmatrix} 3 \\ 5 \\ 9 \end{bmatrix}.$$

The equation $Ax = b$ asks us to express b as a combination of the columns. There is a solution exactly when b is in the **column space**. This leads immediately to the crucial step in understanding linear algebra: to move from two column vectors to the **vector space** of all their combinations.

Here that column space is a plane within three-dimensional space. The student must see it! Not instantly but soon, the imagination fills in the plane of **linear combinations**—the essential operation in linear algebra.

Later come a basis for the column space (from elimination), and its dimension (the rank). The course is built on this mixture of *examples* and *ideas.*

To close this essay, I want to describe four important matrices from applied mathematics. For each of them we ask about pivots and determinants and inverses and eigenvalues. These pages are taken from the very beginning of the new book *Applied Mathematics and Scientific Computing* [3]. I won't include the homework problems...

Thank you for this opportunity to speak about linear algebra.

1.1 Four Special Matrices

An m by n matrix has m rows and n columns and mn entries. We operate on those rows and columns to solve linear systems $Ax = b$ and eigenvalue problems $Ax = \lambda x$. From inputs A and b (and from software like MATLAB) we get outputs x and λ. A fast stable algorithm is extremely important, and this book includes fast algorithms.

One purpose of matrices is to store information, but another viewpoint is equally important for applied mathematics. We see the matrix A as an "operator". *It acts on vectors x to produce Ax.* The components of x have a meaning—displacements or pressures or voltages or prices or concentrations. The operator A also has a meaning—in this chapter it takes differences. Then Ax represents pressure differences or voltage drops or price differentials.

Before we turn the problem over to the machine—and also after, when we interpret the results—it is the meaning we want, as well as the numbers.

This book begins with four special families of matrices—simple and useful, absolutely basic. We look first at the properties of these particular matrices (some properties are obvious and others are hidden). It is terrific to practice linear algebra by working with genuinely important matrices. Here are three in the first family:

$$K_2 = \begin{bmatrix} 2 & -1 \\ -1 & 2 \end{bmatrix} \quad K_3 = \begin{bmatrix} 2 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 2 \end{bmatrix} \quad K_4 = \begin{bmatrix} 2 & -1 & 0 & 0 \\ -1 & 2 & -1 & 0 \\ 0 & -1 & 2 & -1 \\ 0 & 0 & -1 & 2 \end{bmatrix}$$

What is significant about K_2 and K_3 and K_4, and eventually the n by n matrix K_n? I will give six answers in the same order that my class gave them—starting with four properties of the K's that are immediately visible.

1. These matrices are **symmetric**. The entry in row i, column j also appears in row j, column i. Thus $K_{ij} = K_{ji}$, on opposite sides of the main diagonal. Symmetry can be expressed by transposing the whole matrix at once: $K = K^{\mathrm{T}}$.

2. The matrices K_n are **sparse**. Most of their entries are zero when n gets large. K_{1000} has a million entries, but only $1000 + 999 + 999$ are nonzero.

3. The nonzeros lie in a "band" around the main diagonal, so each K_n is *banded.* The band has only three diagonals, so these matrices are **tridiagonal**.

When K is a tridiagonal matrix, the solution of $Ku = f$ is very fast. If the unknown vector u has a thousand components, we can find them in a few thousand steps (which take a small fraction of a second). For a full matrix of order $n = 1000$, solving $Ku = f$ would take many millions of steps. Of course we have to ask if the linear equations have a solution in the first place. That question is coming soon.

4. The matrices have **constant diagonals**. Right away that property wakes up Fourier. It signifies that something is not changing when we move in space or time. The problem is shift-invariant or time-invariant. The matrix is not only tridiagonal, the diagonals are entirely determined by the three numbers $-1, 2, -1$. These are actually "second difference matrices" but my class never says that.

The whole world of Fourier transforms is linked to constant-diagonal matrices. In signal processing, this matrix is a "highpass filter". Ku picks out the rapidly varying (high frequency) part of a vector u. It gives a convolution with $-1, 2, -1$. Mathematicians call K a *Toeplitz matrix*, and MATLAB uses that name:

The command $K = \mathsf{toeplitz}\begin{bmatrix} 2 & -1 & \mathsf{zeros}(1, n-2) \end{bmatrix}$ **constructs** K_n (*first define* n).

We mention these words now to call attention to the Fourier parts of this book. Actually, Fourier will be happier if we make two small changes in K_n. Insert -1 in the southwest and northeast corners. This completes two diagonals (by circling around). All four diagonals wrap around in this "*periodic matrix*" or "*cyclic convolution*" or

circulant matrix:

$$\textbf{Circulant matrix} \qquad C_4 = \begin{bmatrix} 2 & -1 & 0 & \mathbf{-1} \\ -1 & 2 & -1 & 0 \\ 0 & -1 & 2 & -1 \\ \mathbf{-1} & 0 & -1 & 2 \end{bmatrix}.$$

This matrix is *singular*. It is *not invertible*. Its determinant is zero. Rather than computing that determinant, it is much better to identify a nonzero vector u that solves $C_4u = 0$. (**If C_4 had an inverse, the only solution to $C_4u = 0$ would be the zero vector**. We could multiply by C_4^{-1} to find $u = 0$.) For this matrix, the vector $u = (1,1,1,1)$ solves $C_4u = 0$. This all-ones vector is in the *nullspace* of C_4.

Whenever the entries along every row of a matrix add to zero, the matrix is certainly singular. The same all-ones vector u is responsible. Matrix multiplication Ku adds the column vectors and produces zero. This constant vector $u = (1,1,1,1)$ or $u = (c,c,c,c)$ in the nullspace is like the constant C when we integrate in calculus. This "arbitrary constant" is not knowable from the derivative. In linear algebra, the constant in $u = (c,c,c,c)$ is not knowable from $Cu = 0$.

5. All the matrices $K = K_n$ are **invertible**. They are not singular, like C_n. There is a square matrix K^{-1} such that $K^{-1}K = I =$ "*identity matrix*". And if a square matrix has an inverse on the left, then also $KK^{-1} = I$ and the inverse is the same on the right. This "*inverse matrix*" is also symmetric, but K^{-1} is *not* sparse.

Invertibility is not easy to decide from a quick look at a matrix. Theoretically, one test is to compute the determinant of K. There is an inverse except when $\det K = 0$, because the formula for K^{-1} includes a division by $\det K$. But computing the determinant is almost never done in practice!

What we actually do is to go ahead with the elimination steps that solve $Ku = f$. Those steps simplify the matrix, to make it triangular. The nonzero pivots on the main diagonal show that the original K is invertible. (Important: We don't want or need K^{-1} to find $u = K^{-1}f$. The inverse would be a full matrix, with all positive entries. All we compute is the solution vector u.)

6. The symmetric matrices K_n are **positive definite**. One goal of Chapter 1 is to explain what this crucial property means (K_4 has it, C_4 doesn't). Allow me to contrast positive definiteness with invertibility, using words that will soon be familiar. *Please notice the Appendix that summarizes linear algebra.*

(***Pivots***) An invertible matrix has n *nonzero* pivots.
A positive definite symmetric matrix has n *positive* pivots.

(***Eigenvalues***) An invertible matrix has n *nonzero* eigenvalues.
A positive definite symmetric matrix has n *positive* eigenvalues.

Positive pivots and eigenvalues are tests for positive definiteness, and C_4 fails those tests because it is singular. Actually C_4 has three positive pivots and eigenvalues, so it almost passes. But it fails to have a fourth pivot, and its fourth eigenvalue is zero (the matrix is singular). So C_4 is **positive semidefinite**.

The pivots appear on the main diagonal (next section), when solving $Ku = f$. The eigenvalues arise in $Kx = \lambda x$. There is also a determinant test for positive definiteness (not just $\det K > 0$). The proper definition of a symmetric positive definite matrix (it is connected to positive energy) will be linked to these three tests in Section 1.5.

Changing K_n to T_n

After K_n and C_n, there are two more families of matrices that you need to know. They are symmetric and tridiagonal like the family K_n. But the $(1,1)$ entry in T_n is changed from 2 to 1:

$$T_2 = \begin{bmatrix} \mathbf{1} & -1 \\ -1 & 2 \end{bmatrix} \quad \text{and} \quad T_3 = \begin{bmatrix} \mathbf{1} & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 2 \end{bmatrix}. \tag{1}$$

That new top row (T stands for top) represents a new boundary condition, whose meaning we will soon understand. Right now we use T_3 as a perfect example of elimination. Row operations produce zeros below the diagonal, and the pivots are circled. **Two elimination steps reduce T to U.**

Step 1. Add row 1 to row 2, which leaves zeros below the first pivot

Step 2. Add the new row 2 to row 3, which produces an upper triangular U.

$$T = \begin{bmatrix} \textcircled{1} & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 2 \end{bmatrix} \underset{\text{Step 1}}{\longrightarrow} \begin{bmatrix} \textcircled{1} & -1 & 0 \\ 0 & \textcircled{1} & -1 \\ 0 & -1 & 2 \end{bmatrix} \underset{\text{Step 2}}{\longrightarrow} \begin{bmatrix} \textcircled{1} & -1 & 0 \\ 0 & \textcircled{1} & -1 \\ 0 & 0 & \textcircled{1} \end{bmatrix} = U\,.$$

All three pivots of T equal 1. We can apply the test for invertibility (three nonzero pivots) and the test for positive definiteness (three positive pivots). In fact every T_n in this family is positive definite, with its pivots equal to 1.

That upper triangular U has an inverse (U^{-1} is automatically upper triangular). The exceptional fact for this particular U^{-1} is that *all upper triangular entries are* 1's:

$$U^{-1} = \begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & -1 \\ 0 & 0 & 1 \end{bmatrix}^{-1} = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix}. \tag{2}$$

To me, this says that the inverse of a 3 by 3 "difference matrix" is a 3 by 3 "sum matrix". The matrix U takes differences, and U^{-1} takes sums. This neat inverse of U will lead us to the inverse of T. **The product $U^{-1}U$ is the identity matrix I.** Taking differences and then sums will recover the original vector (u_1, u_2, u_3):

$$\textbf{Differences from } \boldsymbol{U}\textbf{:} \quad \begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & -1 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \\ u_3 \end{bmatrix} = \begin{bmatrix} u_1 - u_2 \\ u_2 - u_3 \\ u_3 - 0 \end{bmatrix}$$

$$\textbf{Sums from } \boldsymbol{U^{-1}}\textbf{:} \quad \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} u_1 - u_2 \\ u_2 - u_3 \\ u_3 - 0 \end{bmatrix} = \begin{bmatrix} u_1 \\ u_2 \\ u_3 \end{bmatrix}.$$

Changing T_n to B_n

The fourth family B_n has the last entry also changed from 2 to 1. The new boundary condition is being applied at both ends (B stands for both). These matrices B_n are symmetric and tridiagonal, but you will quickly see that they are *not invertible.* Therefore the B_n are *not positive definite*:

$$B_2 = \begin{bmatrix} \mathbf{1} & -1 \\ -1 & \mathbf{1} \end{bmatrix} \quad \text{and} \quad B_3 = \begin{bmatrix} \mathbf{1} & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & \mathbf{1} \end{bmatrix}. \tag{3}$$

Again, elimination brings out the properties of the matrix. The first $n-1$ pivots will all equal 1, because in those rows there is no change from T_n. But the change from 2 to 1 in the last entry of B produces a change from 1 to 0 in the last entry of U:

$$B = \begin{bmatrix} \textcircled{1} & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & \mathbf{1} \end{bmatrix} \longrightarrow \begin{bmatrix} \textcircled{1} & -1 & 0 \\ 0 & \textcircled{1} & -1 \\ 0 & -1 & \mathbf{1} \end{bmatrix} \longrightarrow \begin{bmatrix} \textcircled{1} & -1 & 0 \\ 0 & \textcircled{1} & -1 \\ 0 & 0 & \mathbf{0} \end{bmatrix} = U\,. \tag{4}$$

There are only two pivots. (A pivot must be nonzero.) The last matrix U is certainly not invertible. Its determinant is zero, because its third row is all zeros. The constant vector $(1,1,1)$ is in the nullspace of U, and therefore it is in the nullspace of B:

$$\begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & -1 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \quad \text{and also} \quad \begin{bmatrix} 1 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}.$$

The whole point of elimination was to simplify a linear system like $Bu = 0$, *without changing the solutions.* In this case we could have recognized non-invertibility in the matrix B, because each row adds to zero. Then the sum of its three columns is the zero column. This is what we see when B multiplies the vector $(1,1,1)$.

Let me summarize this section in four lines (all these matrices are symmetric):

K_n and T_n are invertible and (more than that) positive definite.

C_n and B_n are singular and (more than that) positive semidefinite.

The nullspaces of C_n and B_n contain all the constant vectors $u = (c, c, \ldots, c)$.

The nullspaces of K_n and T_n contain only the zero vector $u = (0, 0, \ldots, 0)$.

Matrices in MATLAB

It is natural to choose MATLAB as the software for linear algebra. This is not the only choice, and the reader may work with another system (Mathematica and Maple are good for symbolic calculation, LAPACK provides excellent codes at no cost, and there are many other linear algebra packages). We will construct matrices and operate on them in the convenient language that MATLAB provides.

Our first step is to construct the matrices K_n. For $n = 3$, we can enter the 3 by 3 matrix a row at a time, inside brackets. Rows are separated by a semicolon ;

$$K = [\ 2 \quad -1 \quad 0\ ; \quad -1 \quad 2 \quad -1\ ; \quad 0 \quad -1 \quad 2\ ]$$

For larger sizes n this will be too slow. Instead we can build K_8 from "eye" and "ones":

$$\mathsf{eye}(8) = 8 \text{ by } 8 \text{ identity matrix} \qquad \mathsf{ones}(7,1) = \text{ column vector of seven 1's}$$

The diagonal part is $2{*}\mathsf{eye}(8)$. Notice the symbol $*$ for multiplication! The -1's above the diagonal of K_8 have the vector $\mathsf{ones}(7,1)$ placed along diagonal 1 of the matrix E:

$$E = -\mathsf{diag}(\mathsf{ones}(7,1),1)$$

The -1's *below* the diagonal of K_8 lie on the diagonal numbered -1. For those we could change the last argument in E from 1 to -1. Or we can simply transpose E, using the all-important symbol E' for E^{T}. Then K comes from its three diagonals:

$$K = 2 * \mathsf{eye}(8) + E + E' \quad \text{is the tridiagonal matrix } K_8.$$

To construct the matrices T and B and C from K, just change entries as in the last three lines of this m-file that we have named special.m. Its input is the size n, its output is four matrices of that size. The semicolons stop display of K, T, B, C:

```
function [K,T,B,C] = special(n)

K = toeplitz ([2 -1 zeros(1,n-2)]);
T = K; T(1,1) = 1;
B = K; B(1,1) = 1; B(n,n) = 1;
C = K; C(1,n) = -1; C(n,1) = -1;
```

If we happened to want their determinants (we shouldn't!), then with $n = 8$

$$[\ \det(K) \quad \det(T) \quad \det(B) \quad \det(C)\ ] \text{ produces the output} \qquad 9 \quad 1 \quad 0 \quad 0$$

One more important point. MATLAB could not store K_n as a dense matrix for $n = 10{,}000$. **The 10^8 entries need 800MB unless we recognize K as sparse**. The code KTBC.m on the course website avoids storing (and operating on) all the zeros. It has K, T, B, or C as first argument and n as second argument. The third argument is 1 for sparse, 0 for dense (default 0 for $\mathsf{narg} = 2$, no third argument).

The input to Sparse MATLAB includes the locations of all nonzero entries. Our KTBC code uses spdiags to enter the diagonals. The more general command $A = \mathsf{sparse}(i, j, s, m, n)$ creates an m by n sparse matrix from the vectors i, j, s that list positions i, j of nonzero entries s. Elimination by $\mathsf{lu}(A)$ may produce additional nonzeros (called fill-in) which the software will correctly identify. In the normal "dense" option, zeros are processed like all other numbers.

The next sections will use all four matrices in these basic tasks of linear algebra:

(**1.2**) Elimination produces pivots in D and triangular factors in LDL^{T}

(**1.3**) The finite difference matrices K, T, B, C come with boundary conditions

(**1.4**) Point loads produce inverse matrices K^{-1} and T^{-1}

(**1.5**) The eigenvalues and eigenvectors involve sines and cosines.

I very much hope that you will come to know and like these special matrices.

References (with apologies)

[1] Gilbert Strang, *Introduction to Linear Algebra*, Wellesley-Cambridge Press (Third edition 2003).

[2] Gilbert Strang, *Linear Algebra and Its Applications*, Brooks/Cole (Third edition 1988, Fourth edition 2005).

[3] Gilbert Strang, *Applied Mathematics and Scientific Computing*, Wellesley-Cambridge Press (2005).