A Tour of Linear Algebra for Math 4/896, Sec. 006

Note to 496 students: This really is a *brief* tour. You will find much more detail in the excellent linear algebra review Appendix A of our textbook. Everyone should read through this appendix. For the most part, it is material from Math 314. Since it is less likely that you saw a variety of norms and a development of the singular value decomposition, I'm including a little more detail concerning these topics.

Solving Systems

Solving linear systems is a fundamental activity, practiced from high school forward. Recall that a linear equation is one of the form

$$a_1x_1 + a_2x_2 + \cdots + a_nx_n = b,$$

where $x_1, x_2, \ldots, x_n$ are the variables and $a_1, a_2, \ldots, a_n, b$ are known constants. A linear system is a collection of linear equations, and a solution to the linear system is a choice of the variables that satisfies every equation in the system. What can we expect?

Example 6. Use geometry to answer the question with one or more equations in 2 unknowns. Ditto 3 unknowns.

Basic Fact (Number of Solutions): A linear system of m equations in n unknowns is either inconsistent (no solutions) or consistent, in which case it either has a unique solution or infinitely many solutions. In the former case, unknowns cannot exceed equations. If the number of equations exceeds the number of unknowns, the system is called **overdetermined**. If the number of unknowns exceeds the number of equations, the system is called **underdetermined**.

While we're on this topic, let's review some basic linear algebra as well:

The general form for a linear system of m equations in the n unknowns $x_1, x_2, \ldots, x_n$ is

$$\begin{aligned}
a_{11}x_1 + a_{12}x_2 + \cdots + a_{1j}x_j + \cdots a_{1n}x_n &= b_1 \\
a_{21}x_1 + a_{22}x_2 + \cdots + a_{2j}x_j + \cdots a_{2n}x_n &= b_2 \\
\vdots \quad &\vdots \quad \vdots \\
a_{i1}x_1 + a_{i2}x_2 + \cdots + a_{ij}x_j + \cdots a_{in}x_n &= b_i \\
\vdots \quad &\vdots \quad \vdots \\
a_{m1}x_1 + a_{m2}x_2 + \cdots + a_{mj}x_j + \cdots a_{mn}x_n &= b_m
\end{aligned}$$

Notice how the coefficients are indexed: in the ith row the coefficient of the jth variable, x_j, is the number a_{ij}, and the right hand side of the ith equation is b_i.

The statement "$A = [a_{ij}]$" means that A is a **matrix** (rectangular table of numbers) whose (i, j)th entry, i.e., entry in the ith row and jth column, is denoted by a_{ij}. Generally, the size of A will be clear from context. If we want to indicate that A is an $m \times n$ matrix, we write

$$A = [a_{ij}]_{m,n}.$$

Similarly, the statement "$\mathbf{b} = [b_i]$" means that b is a **column vector** (matrix with exactly one column) whose ith entry is denoted by b_i, and "$\mathbf{c} = [c_j]$" means that c is a **row vector** (matrix with exactly one row) whose jth entry is denoted by c_j. In case the type of the vector (row or column) is not clear from context, the default is a column vector.

How can we describe the matrices of the general linear system described above? First, there is the $m \times n$ **coefficient matrix**

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1j} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2j} & \cdots & a_{2n} \\ \vdots & \vdots & & \vdots & & \vdots \\ a_{i1} & a_{i2} & \cdots & a_{ij} & \cdots & a_{in} \\ \vdots & \vdots & & \vdots & & \vdots \\ a_{m1} & a_{m2} & \cdots & a_{mj} & \cdots & a_{mn} \end{bmatrix}$$

Notice that the way we subscripted entries of this matrix is really very descriptive: the first index indicates the row position of the entry and the second, the column position of the entry. Next, there is the $m \times 1$ **right hand side vector** of constants

$$\mathbf{b} = \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_i \\ \vdots \\ b_m \end{bmatrix}$$

Finally, stack this matrix and vector along side each other (we use a vertical bar below to separate the two symbols) to obtain the $m \times (n+1)$

augmented matrix

$$\widetilde{A} = [A \mid \mathbf{b}] = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1j} & \cdots & a_{1n} & b_1 \\ a_{21} & a_{22} & \cdots & a_{2j} & \cdots & a_{2n} & b_2 \\ \vdots & \vdots & & \vdots & & \vdots & \vdots \\ a_{i1} & a_{i2} & \cdots & a_{ij} & \cdots & a_{in} & b_i \\ \vdots & \vdots & & \vdots & & \vdots & \vdots \\ a_{m1} & a_{m2} & \cdots & a_{mj} & \cdots & a_{mn} & b_m \end{bmatrix}$$

Matrix Arithmetic

Given any two matrices (or vectors) $A = [a_{ij}]$ and $B = [b_{ij}]$ of the same size, we may add them or multiply by a scalar according to the rules

$$A + B = [a_{ij} + b_{ij}]$$

and

$$cA = [ca_{ij}] .$$

For each size $m \times n$ we can form a **zero matrix** of size $m \times n$: just use zero for every entry. This is a handy matrix for matrix arithmetic. Matlab knows all about these arithmetic operations, as we saw in our Matlab introduction. Moreover, there is a multiplication of matrices A and B provided that A is $m \times p$ and B is $p \times n$. The result is an $m \times n$ matrix given by the formula

$$AB = \left[\sum_{k=1}^{p} a_{ik} b_{kj} \right] .$$

Again, Matlab knows all about matrix multiplication.

The Basics. There are many algebra rules for matrix arithmetic. We won't review all of them here, but one noteworthy fact is that matrix multiplication is not commutative: $AB \neq BA$ in general. Another is that there is no general cancellation of matrices: if $AB = 0$. Another is that there is an **identity matrix** for any square size $n \times n$: I_n is the matrix with ones down the diagonal entries and zeros elsewhere. For example,

$$I_3 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} .$$

We normally just write I and n should be clear from context. Algebra fact: if A is $m \times n$ and B is $n \times m$, then $AI = A$ and $IB = B$.

One of the main virtues of matrix multiplication is that is gives us a symbolic way of representing the general linear system given above, namely, we can put it in the simple looking form

$$A\mathbf{x} = \mathbf{b},$$

where A and b are given as above, namely

$$A\mathbf{x} = \mathbf{b},$$

where

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_i \\ \vdots \\ x_n \end{bmatrix}$$

is the (column) vector of unknowns. Actually, this is putting the cart before the horse because this idea actually inspires the whole notion of matrix multiplication. This form suggests a way of thinking about the solution to the system, namely,

$$\mathbf{x} = \frac{\mathbf{b}}{A}.$$

The only problem is that this doesn't quite make sense in the matrix world. But we can make it sensible by a little matrix algebra: imagine that there were a multiplicative inverse matrix A^{-1} for A, that is, a matrix such that

$$A^{-1}A = I = AA^{-1}.$$

Then we could multiply both sides of $Ax = b$ on the left by A^{-1} and obtain that

$$\mathbf{x} = A^{-1}\mathbf{b}.$$

Matlab represents this equation by
x = A\b

In general, a square matrix A may not have an inverse. One condition for an inverse to exist is that the matrix have nonzero determinant, where this is a number associated with a matrix such as one saw in high school algebra with Cramer's rule. We won't go into details here. Of course, Matlab knows all about inverses and we can compute the inverse of A by issuing the commands A^(-1) or inv(A).

Sets of *all* matrices or vectors of a given size form "number systems" with a scalar multiplication and addition operation that satisfies a list of basic laws. Such a system is a **vector space.** For example, for vectors from $\mathbb{R}^3$, we have operations

$$\begin{bmatrix} a_1 \\ a_2 \\ a_3 \end{bmatrix} + \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix} = \begin{bmatrix} a_1 + b_1 \\ a_2 + b_2 \\ a_3 + b_3 \end{bmatrix}$$

$$c \begin{bmatrix} a_1 \\ a_2 \\ a_3 \end{bmatrix} = \begin{bmatrix} ca_1 \\ ca_2 \\ ca_3 \end{bmatrix}$$

and similar operations for n-vectors of the vector space $\mathbb{R}^n$ or the vector space of $n \times n$ matrices $\mathbb{R}^{n \times n}$.

A more exotic example of a vector space (I'm not bothering to give the formal definition) is

$$C[a,b] = \{f(x) : [a,b] \to \mathbb{R} \mid f(x) \text{ is continuous on } [a,b]\}$$

together with the vector space operations of addition and scalar multiplication given by the formulas

$$\begin{aligned} (f+g)(x) &= f(x) + g(x) \\ (cf)(x) &= c(f(x)). \end{aligned}$$

With vector spaces come a steady flow of ideas that we won't elaborate on here – such as subspaces, linear independence, spanning sets, bases and many more.

Inner Products and Norms

Norms. Before we get started, there are some useful ideas that we need to explore. We already know how to measure the size of a scalar (number) quantity x: use $|x|$ as a measure of its size. Thus, we have a way of thinking about things like the *size* of an error. Now suppose we are dealing with vectors and matrices.

Question: How do we measure the size of more complex quantities like vectors or matrices?

Answer: We use some kind of yardstick called a **norm** that assigns to each vector $\mathbf{x}$ a non-negative number $\|\mathbf{x}\|$ subject to the following norm laws for arbitrary vectors $\mathbf{x}, \mathbf{y}$ and scalar c:

- For $\mathbf{x} \neq \mathbf{0}$, $\|\mathbf{x}\| > 0$ and for $\mathbf{x} = \mathbf{0}$, $\|\mathbf{x}\| = 0$.
- $\|c\mathbf{x}\| = |c| \, \|\mathbf{x}\|$.
- $\|\mathbf{x} + \mathbf{y}\| \leq \|\mathbf{x}\| + \|\mathbf{y}\|$.

Examples: For vector space $\mathbb{R}^n$ and vector $\mathbf{x} = [x_1, x_2; \dots; x_n]$,

- 1-norm: $\|\mathbf{x}\|_1 = |x_1| + |x_2| + \cdots + |x_n|$

- 2-norm (the default norm):

$$\|\mathbf{x}\|_2 = \left(x_1^2 + x_2^2 + \cdots + x_n^2\right)^{1/2}$$

- ∞-norm:$\|\mathbf{x}\|_\infty = \max\left(|x_1|, |x_2|, \cdots, |x_n|\right)$

Let's do some examples in Matlab, which knows all of these norms:

```
  >A=[1;-3;2;-1], x=[2;0;-1;2]
>norm(x)
>norm(x,1)
>norm(x,2)
>norm(x,inf)>norm(x)
>norm(A,1)
>norm(A,2)
>norm(A,inf)
>norm(-2*x),abs(-2)*norm(x)
>norm(A*x),norm(A)*norm(x)
```

We can abstract the norm idea to more exotic vector spaces like $C[a,b]$ in many ways. Here is one that is analogous to the infinity norm on $\mathbb{R}^n$:

$$\|f\| \equiv \max_{a\le x\le b} |f(x)|.$$

Matrix Norms: For every vector norm $\|\mathbf{x}\|$ of n-vectors there is a corresponding induced norm on the vector space of $n \times n$ matrices A by the formula

$$\|A\| = \max_{\mathbf{x}\neq\mathbf{0}} \frac{\|A\mathbf{x}\|}{\|\mathbf{x}\|} = \max_{\|\mathbf{x}\|=1} \|A\mathbf{x}\|.$$

These induced norms satisfy all the basic norm laws and, in addition:

$$\begin{aligned} \|A\mathbf{x}\| &\le \|A\|\,\|\mathbf{x}\| \text{ (compatible norms)} \\ \|AB\| &\le \|A\|\,\|B\| \text{ (multiplicative norm)} \end{aligned}$$

One can show directly that if $A = [a_{ij}]$, then

- $\|A\|_1 = \max\limits_{1\le j\le n} \sum\limits_{i=1}^{n} |a_{ij}|$
- $\|A\|_2 = \sqrt{\rho\left(A^T A\right)}$, where $\rho(B)$ is largest eigenvalue of square matrix B in absolute value.
- $\|A\|_\infty = \max\limits_{1\le i\le n} \sum\limits_{j=1}^{n} |a_{ij}|$

There is a more general definition of matrix norm that requires that the basic norm laws be satisfied as well as one more, the so-called multiplicative property:

$$\|AB\| \leq \|A\| \, \|B\|$$

for all matrices conformable for multiplication. One such standard example is the **Frobenius norm**: given that $A = [a_{ij}]$

$$\|A\|_F = \sum_{i=1}^{m} \sum_{j=1}^{n} |a_{ij}|^2 \, .$$

One can check that the all the matrix norm properties are satisfied.

Inner Products. The 2-norm is special. It can be derived from another operation on vectors called the inner product (or dot product). Here is how it works.

Definition. The **inner product** (dot product) of vectors $\mathbf{u}, \mathbf{v}$ in $\mathbb{R}^n$ is the scalar

$$\mathbf{u} \cdot \mathbf{v} = \mathbf{u}^T \mathbf{v}.$$

This operation satisfies a long list of properties, called the **inner product properties**:

- $\mathbf{u} \cdot \mathbf{u} \geq 0$ with equality if and only if $\mathbf{u} = \mathbf{0}$.
- $\mathbf{u} \cdot \mathbf{v} = \mathbf{v} \cdot \mathbf{u}$
- $\mathbf{u} \cdot (\mathbf{v} + \mathbf{w}) = \mathbf{u} \cdot \mathbf{v} + \mathbf{u} \cdot \mathbf{w}$
- $c \, (\mathbf{u} \cdot \mathbf{v}) = (c\mathbf{u}) \cdot \mathbf{v}$

What really makes this idea important are the following facts. Here we understand that $\|\mathbf{u}\|$ means the 2-norm $\|\mathbf{u}\|_2$.

- $$\|\mathbf{u}\| = \sqrt{\mathbf{u} \cdot \mathbf{u}}$$
- If θ is an angle between the vectors represented in suitable coordinate space, then
  $$\mathbf{u} \cdot \mathbf{v} = \|\mathbf{u}\| \, \|\mathbf{u}\| \cos \theta.$$
- If θ is an angle between the vectors represented in suitable coordinate space, then
  $$\mathbf{u} \cdot \mathbf{v} = \|\mathbf{u}\| \, \|\mathbf{u}\| \cos \theta.$$
- Once we have the idea of angle between vectors, we can speak of **orthogonal** (perpendicular) vectors, i.e., vectors for which the angle between them is $\pi/2$ (90°).

- Given any two vectors $\mathbf{u}, \mathbf{v}$, with $\mathbf{v} \neq \mathbf{0}$, the formula
$$\operatorname{proj}_{\mathbf{v}} \mathbf{u} = \frac{\mathbf{u} \cdot \mathbf{v}}{\mathbf{v} \cdot \mathbf{v}} \mathbf{v}$$
defines a vector parallel to $\mathbf{v}$ such that $\mathbf{u} - \operatorname{proj}_{\mathbf{v}} \mathbf{u}$ is orthogonal to $\mathbf{v}$.
- We can introduce the concept of **orthogonal set of vectors**: An orthogonal set of vectors is one such that elements are pairwise orthogonal. It turns out that orthogonal sets of nonzero vectors are always linearly independent. The Gram-Schmidt algorithm gives us a way of turning any linearly independent set into an orthogonal set with the same span.

We can abstract the idea of inner product to more exotic vector spaces. Consider, for example, $C[a, b]$ with the inner product (which we now denote with angle brackets instead of the dot which is reserved for the standard vector spaces)
$$\langle f, g \rangle = \int_a^b f(x)\, g(x)\, dx.$$
Of course, this gives rise to an induced norm
$$\|f\| = \int_a^b f(x)^2\, dx$$

Eigenvectors and Eigenvalues

An **eigenvalue** for a square $n \times n$ matrix A is a number λ such that for some NONZERO vector $\mathbf{x}$, called an **eigenvector** for λ,
$$A\mathbf{x} = \lambda \mathbf{x}.$$

Eigenvalue Facts:

(1) The eigenvalues of A are precisely the roots of the nth degree polynomial
$$p(\lambda) = |\lambda I - A|,$$
where $||$ denotes determinant.

(2) An $n \times n$ matrix A has exactly n eigenvalues, counting multiplicities and complex numbers.

(3) The eigenvalues of $\alpha I + \beta A$ are just $\alpha + \beta\lambda$, where λ runs over the eigenvalues of A.

(4) More generally, if $r(\lambda) = p(\lambda)/q(\lambda)$ is any rational function, i.e., p, q are polynomials, then the eigenvalues of $r(A)$ are $r(\lambda)$, where λ runs over the eigenvalues of A, provided both expressions make sense.

(5) The **spectral radius** of A is just

$$\rho(A) = \max\{|\lambda| \mid \lambda \text{ is an eigenvalue of } A\}$$

(6) Suppose that $\mathbf{x}_0$ is an arbitrary initial vector, $\{\mathbf{b}_k\}$ is an arbitrary sequence of uniformly bounded vectors, and the sequence $\{\mathbf{x}_k\}$, $k = 0, 1, 2, \ldots$, is given by the formula

$$\mathbf{x}_{k+1} = A\mathbf{x}_k + \mathbf{b}_k.$$

If $\rho(A) < 1$, then $\{\mathbf{x}_k\}$ is a uniformly bounded sequence of vectors. Such a matrix A is called a **stable** matrix.

(7) If $\rho(A) > 1$, then $\{\mathbf{x}_k\}$ will not be uniformly bounded for some choices of $\mathbf{x}_0$.

Let's do a few calculations with matrices to illustrate the idea of eigenvalues and eigenvectors in Matlab, which knows all about these quantities.

```
> A = [3 1 0;-1 3 1;0 3 -2]
> eig(A)
> eig(A)
> [V,D]=eig(A)
> v = V(:,1),lam = D(1,1)
> A*v,lam*v
> x = [1;-2;3]
> b = [2;1;0]
> x = A*x+b % repeat this line
```

Now let's demonstrate boundedness of an arbitrary sequence as above. Try the same with a suitable A.

```
> A = A/4
> max(abs(eig(A)))
> x = [1;-2;3]
> x = A*x+b % repeat this line
```

Condition Number of a Matrix

Condition number is only defined for invertible square matrices.

Definition. Let A be a square invertible matrix. The **condition number** of A with respect to a matrix norm $\|\cdot\|$is the number

$$\operatorname{cond}(A) = \|A\| \left\|A^{-1}\right\|.$$

Roughly speaking, the condition number measures the degree to which changes in A lead to changes in solutions of systems $A\mathbf{x} = \mathbf{b}$. A large condition number means that small changes in A may lead to large changes in $\mathbf{x}$. Of course this quantity is norm dependent.

Definition. One of the more important applications of the idea of a matrix norm is the famous Banach Lemma. Essentially, it amounts to a matrix version of the familiar geometric series encountered in calculus.

Theorem 0.1. *Let M be a square matrix such that $\|M\| < 1$ for some operator norm $\|\cdot\|$. Then the matrix $I - M$ is invertible. Moreover,*

$$(I - M)^{-1} = I + M + M^2 + \cdots + M^k + \cdots$$

and $\left\|(I - M)^{-1}\right\| \leq 1/(1 - \|M\|)$.

Proof. Form the familiar telescoping series

$$(I - M)\left(I + M + M^2 + \cdots + M^k\right) = I - M^{k+1}$$

so that

$$I - (I - M)\left(I + M + M^2 + \cdots + M^k\right) = M^{k+1}$$

Now by the multiplicative property of matrix norms and fact that $\|M\| < 1$

$$\left\|M^{k+1}\right\| \leq \|M\|^{k+1} \underset{k\to\infty}{\to} 0.$$

It follows that the matrix $\lim_{k\to\infty}\left(I + M + M^2 + \cdots + M^k\right) = B$ exists and that $I - (I - M)B = 0$, from which it follows that $B = (I - M)^{-1}$. Finally, note that

$$\begin{aligned}\left\|I + M + M^2 + \cdots + M^k\right\| &\leq \|I\| + \|M\| + \|M\|^2 + \cdots + \|M\|^k \\ &\leq 1 + \|M\| + \|M\|^2 + \cdots + \|M\|^k \\ &\leq \frac{1}{1 - \|M\|}.\end{aligned}$$

Now take the limit as $k \to \infty$ to obtain the desired result. □

In the case of an operator norm, the Banach lemma has a nice application.

Corollary 0.2. *If $A = I + N$, where $\|N\| < 1$, then*

$$\mathrm{cond}(A) \leq \frac{1 + \|N\|}{1 - \|N\|}$$

We leave the proof as an exercise.

Here is a very fundamental result for numerical linear algebra. The scenario: suppose that we desire to solve the linear system $A\mathbf{x} = \mathbf{b}$, where A is invertible. Due to arithmetic error or possibly input data error, we end up with a value $\mathbf{x} + \delta\mathbf{x}$ which solves exactly a "nearby" system $(A + \delta A)(\mathbf{x} + \delta\mathbf{x}) = \mathbf{b} + \delta\mathbf{b}$. (It can be shown by using an idea called "backward error analysis" that this is really what happens when many algorithms are used to solve a linear system.) The question is, what is the size of the relative error $\|\delta x\|/\|x\|$? As long as the perturbation matrix $\|\delta A\|$ is reasonably small, there is a very elegant answer.

Theorem 0.3. *Suppose that A is invertible, $A\mathbf{x} = \mathbf{b}$, $(A + \delta A)(\mathbf{x} + \delta\mathbf{x}) = \mathbf{b} + \delta\mathbf{b}$ and $\|A^{-1}\delta A\| = c < 1$ with respect to some operator norm. Then $A + \delta A$ is invertible and*

$$\frac{|\ \delta\mathbf{x}\|}{\|\mathbf{x}\|} \leq \frac{\mathrm{cond}(A)}{1-c}\left[\frac{\|\delta A\|}{\|A\|} + \frac{\|\delta\mathbf{b}\|}{\|\mathbf{b}\|}\right]$$

Proof. That the matrix $I + A^{-1}\delta A$ follows from hypothesis and the Banach lemma. Expand the perturbed equation to obtain

$$(A + \delta A)(\mathbf{x} + \delta\mathbf{x}) = A\mathbf{x} + \delta A\mathbf{x} + A\delta\mathbf{x} + \delta A\,\delta\mathbf{x} = \mathbf{b} + \delta\mathbf{b}$$

Now subtract the terms $A\mathbf{x} = \mathbf{b}$ from each side and solve for $\delta\mathbf{x}$ to obtain

$$(A + \delta A)\delta\mathbf{x} = A^{-1}(I + A^{-1}\delta A)\delta\mathbf{x} = -\delta A \cdot \mathbf{x} + \delta\mathbf{b}$$

so that

$$\delta\mathbf{x} = (I + A^{-1}\delta A)^{-1}A^{-1}\left[-\delta A \cdot \mathbf{x} + \delta\mathbf{b}\right]$$

Now take norms and use the additive and multiplicative properties and the Banach lemma to obtain

$$\|\delta\mathbf{x}\| \leq \frac{\|A^{-1}\|}{1-c}\left[\|\delta A\|\|\mathbf{x}\| + \|\delta\mathbf{b}\|\right].$$

Next divide both sides by $\|x\|$ to obtain

$$\frac{\|\delta\mathbf{x}\|}{\|\mathbf{x}\|} \leq \frac{\|A^{-1}\|}{1-c}\left[\|\delta A\| + \frac{\|\delta\mathbf{b}\|}{\|\mathbf{x}\|}\right]$$

Finally, notice that $\|\mathbf{b}\| \leq \|A\|\|x\|$. Therefore, $1/\|\mathbf{x}\| \leq \|A\|/\|\mathbf{b}\|$. Replace $1/\|\mathbf{x}\|$ in the right hand side by $\|A\|/\|\mathbf{b}\|$ and factor out $\|A\|$ to obtain

$$\frac{\|\delta\mathbf{x}\|}{\|\mathbf{x}\|} \leq \frac{\|A^{-1}\|\|A\|}{1-c}\left[\frac{\|\delta A\|}{\|A\|} + \frac{\|\delta\mathbf{b}\|}{\|\mathbf{b}\|}\right]$$

which completes the proof, since by definition, $\mathrm{cond}\,A = \|A^{-1}\|\|A\|$. □

If we believe that the inequality in the perturbation theorem can be sharp (it can!), then it becomes clear how the condition number of the matrix A is a direct factor in how relative error in the solution vector is amplified by perturbations in the coefficient matrix.

Eigenvectors and Eigenvalues

An **eigenvalue** for a square $n \times n$ matrix A is a number λ such that for some NONZERO vector $\mathbf{x}$, called an **eigenvector** for λ,

$$A\mathbf{x} = \lambda\mathbf{x}.$$

Eigenvalue Facts:

(1) The eigenvalues of A are precisely the roots of the nth degree polynomial

$$p(\lambda) = |\lambda I - A|,$$

where $||$ denotes determinant.

(2) An $n \times n$ matrix A has exactly n eigenvalues, counting multiplicities and complex numbers.

(3) The eigenvalues of $\alpha I + \beta A$ are just $\alpha + \beta\lambda$, where λ runs over the eigenvalues of A.

(4) More generally, if $r(\lambda) = p(\lambda)/q(\lambda)$ is any rational function, i.e., p, q are polynomials, then the eigenvalues of $r(A)$ are $r(\lambda)$, where λ runs over the eigenvalues of A, provided both expressions make sense.

(5) The **spectral radius** of A is just

$$\rho(A) = \max\{|\lambda| \mid \lambda \text{ is an eigenvalue of } A\}$$

(6) Suppose that $\mathbf{x}_0$ is an arbitrary initial vector, $\{\mathbf{b}_k\}$ is an arbitrary sequence of uniformly bounded vectors, and the sequence $\{\mathbf{x}_k\}$, $k = 0, 1, 2, \ldots$, is given by the formula

$$\mathbf{x}_{k+1} = A\mathbf{x}_k + \mathbf{b}_k.$$

If $\rho(A) < 1$, then $\{\mathbf{x}_k\}$ is a uniformly bounded sequence of vectors. Such a matrix A is called a **stable** matrix.

(7) If $\rho(A) > 1$, then $\{\mathbf{x}_k\}$ will not be uniformly bounded for some choices of $\mathbf{x}_0$.

(8) The matrix A is diagonalizable, i.e., there exists an invertible matrix P and diagonal matrix D such that

$$P^{-1}AP = D = \begin{bmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \ddots & \\ & & & \lambda_n \end{bmatrix}.$$

Here the diagonal entries of D are precisely the eigenvalues of A.

Let's do a few calculations with matrices to illustrate the idea of eigenvalues and eigenvectors in Matlab, which knows all about these quantities.

```
  > A = [3 1 0;-1 3 1;0 3 -2]
> eig(A)
> eig(A)
> [V,D]=eig(A)
> v = V(:,1),lam = D(1,1)
> A*v,lam*v
> x = [1;-2;3]
> b = [2;1;0]
> x = A*x+b % repeat this line
  Now let's demonstrate boundedness of an arbitrary sequence
as above.  Try the same with a suitable A.
  > A = A/4
> max(abs(eig(A)))
> x = [1;-2;3]
> x = A*x+b % repeat this line
```

Singular Value Decomposition

Here is the fundamental result:

Theorem 0.4. *Let A be an $m \times n$ real matrix. Then there exist $m \times m$ orthogonal matrix U, $n \times n$ orthogonal matrix V and $m \times n$ diagonal matrix Σ with diagonal entries $\sigma_1 \geq \sigma_2 \geq \ldots \geq \sigma_p$, with $p = \min\{m, n\}$, such that $U^T AV = \Sigma$. Moreover, the numbers $\sigma_1, \sigma_2, \ldots, \sigma_p$ are uniquely determined by A.*

Proof. There is no loss of generality in assuming that $n = \min\{m, n\}$. For if this is not the case, we can prove the theorem for A^T and by transposing the resulting SVD for A^T, obtain a factorization for A. Form the $n \times n$ matrix $B = A^T A$. This matrix is symmetric and its eigenvalues are nonnegative (we leave these facts as exercises). Because they are nonnegative, we can write the eigenvalues of B in decreasing order of magnitude as the squares of positive real numbers, say as $\sigma_1^2 \geq \sigma_2^2 \geq \ldots \geq \sigma_n^2$. Now we know from the Principal Axes Theorem that we can find an orthonormal set of eigenvectors corresponding to these eigenvalues, say $B\mathbf{v}_k = \sigma_k^2 \mathbf{v}_k, \quad k = 1, 2, \ldots, n$. Let $V = [\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_n]$. Then V is an orthogonal $n \times n$ matrix. Next, suppose that $\sigma_{r+1}, \sigma_{r+2}, \ldots, \sigma_n$ are zero, while $\sigma_r \neq 0$.

Next set $\mathbf{u}_j = \frac{1}{\sigma_j} A\mathbf{v}_j,\ j = 1, 2, \ldots, r.$ These are orthonormal vectors in $\mathbb{R}^m$ since

$$\mathbf{u}_j^T \mathbf{u}_k = \frac{1}{\sigma_j \sigma_k} \mathbf{v}_j^T A^T A \mathbf{v}_k = \frac{1}{\sigma_j \sigma_k} \mathbf{v}_j^T B \mathbf{v}_k = \frac{\sigma_k^2}{\sigma_j \sigma_k} \mathbf{v}_j^T \mathbf{v}_k = \begin{cases} 0, & \text{if } j \neq k \\ 1, & \text{if } j = k \end{cases}$$

Now expand this set to an orthonormal basis $\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_m$ of $\mathbb{R}^m$. This is possible by standard theorems in linear algebra. Now set $U = [\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_m]$. This matrix is orthogonal and we calculate that if $k > r$, then $\mathbf{u}_j^T A\mathbf{v}_k = 0$ since $A\mathbf{v}_k = 0$, and if $k < r$, then

$$\mathbf{u}_j^T A\mathbf{v}_k = \mathbf{u}_j^T A\mathbf{v}_k = \sigma_k \mathbf{u}_j^T \mathbf{u}_k = \begin{cases} 0, & \text{if } j \neq k \\ \sigma_k, & \text{if } j = k \end{cases}$$

It follows that $U^T AV = [\mathbf{u}_j^T A\mathbf{v}_k] = \Sigma$.

Finally, if U, V are orthogonal matrices such that $U^T AV = \Sigma$, then $A = U\Sigma V^T$ and therefore

$$B = A^T A = V\Sigma U^T U\Sigma V^T = V\Sigma^2 V^T$$

so that the squares of the diagonal entries of Σ are the eigenvalues of B. It follows that the numbers $\sigma_1, \sigma_2, \ldots, \sigma_n$ are uniquely determined by A. □

Notation 0.5. The numbers $\sigma_1, \sigma_2, \ldots, \sigma_p$ are called the **singular values** of the matrix A, the columns of U are the **left singular vectors** of A, and the columns of V are the **right singular values** of A.

There is an interesting geometrical interpretation of this theorem from the perspective of linear transformations and change of basis as developed in Section **??**. It can can be stated as follows.

Corollary 0.6. *Let $T : \mathbb{R}^n \to \mathbb{R}^m$ be a linear transformation with matrix A with respect to the standard bases. Then there exist orthonormal bases $\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_m$ and $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_n$ of $\mathbb{R}^m$ and $\mathbb{R}^n$, respectively, such that the matrix of T with these bases is diagonal with nonnegative entries down the diagonal.*

Proof. First observe that if $U = [\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_m]$ and $V = [\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_n]$, then U and V are the change of basis matrices from the standard bases to the bases $\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_m$ and $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_n$ of $\mathbb{R}^m$ and $\mathbb{R}^n$, respectively. Also, $U^{-1} = U^T$. Now apply a change of basis theorem and the result follows. □

We leave the following fact as an exercise.

Corollary 0.7. *Let $U^T AV = \Sigma$ be the SVD of A and suppose that $\sigma_r \neq 0$ and $\sigma_{r+1} = 0$. Then*

(1) $\operatorname{rank} A = r$
(2) $N(A) = \operatorname{span}\{\mathbf{v}_{r+1}, \mathbf{v}_{r+2}, \ldots, \mathbf{v}_n\}$
(3) $\operatorname{range} A = \operatorname{span}\{\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_r\}$

We have only scratched the surface of the many facets of the SVD. Like most good ideas, it is rich in applications. We mention one more. It is based on the following fact, which can be proved by examining the entries of $A = U\Sigma V^T$: The matrix A of rank r can be written as a sum of r rank one matrices, namely

$$A = \sigma_1 \mathbf{u}_1 \mathbf{v}_1^T + \sigma_2 \mathbf{u}_2 \mathbf{v}_2^T + \cdots + \sigma_r \mathbf{u}_r \mathbf{v}_r^T$$

where the $\sigma_k, \mathbf{u}_k, \mathbf{v}_k$ are the singular values, left and right singular vectors, respectively. In fact, it can be shown that this representation is the most economical in the sense that the partial sums

$$A_k = \sigma_1 \mathbf{u}_1 \mathbf{v}_1^T + \sigma_2 \mathbf{u}_2 \mathbf{v}_2^T + \cdots + \sigma_k \mathbf{u}_k \mathbf{v}_k^T, \qquad k = 1, 2, \ldots, r$$

give the rank k approximation to A that is closest among all rank k approximations to A. Thus, if A has a small number of relatively large singular values, say k of them, then we can approximate A pretty well by A_k.