Lecture 2: Matrix Analysis and Norms

1. Orthogonal Matrix Factorization and Applications.

Amid the millennial fever that swept a decade ago, a group of numerical analysts compiled a list of the 20th Century's 'Top 10 Algorithms' (see *Computing in Science & Engineering* vol. 2 (2000)). Prominent among these is 'the decompositional approach to matrix computations', a clunky name but an essential idea that has enabled great progress in numerical analysis over the past 50 years.

Before this breakthrough, typical matrix problems were solved by focusing on manipulations made to the individual entries of a matrix, a recipe for tedium. In place of this technical approach, algorithms are now formulated by first *decomposing* a generic matrix into the product of several simpler matrices, each of which is much easier to work with than the original matrix. This new perspective facilitates the design and analysis of algorithms; at some point one must inevitably handle individual elements, but this should not be the initial focus.

We begin this course by studying one of these decompositions, the *QR factorization*—a matrix $\mathbf{A}$ can be written as the product of a unitary matrix and an upper triangular matrix. This tool will reappear throughout the semester as we solve linear systems, least squares problems, and eigenvalue problems. We shall also study other matrix factorizations, including the singular value decomposition (SVD), the LU decomposition, and the Schur decomposition.

Before we embark on our study of such algorithms, we must review a few basic concepts from matrix theory and establish a method for measuring the size of vectors and matrices.

1.1. Concepts from matrix analysis.

Throughout these notes, matrices are denoted by bold capital letters; column vectors are denoted with bold lower case letters; scalars will never be bold, and almost always will be lower case. The matrix $\mathbf{A}$ with m rows and n columns with real entries is denoted $\mathbf{A} \in \mathbb{R}^{m\times n}$; if $\mathbf{A}$ has complex entries, we write $\mathbf{A} \in \mathbb{C}^{m\times n}$. The element of $\mathbf{A}$ in row j and column k is denoted a_{jk}.

A set of vectors $\mathcal{U} \subset \mathbb{C}^n$ is a *subspace* if it is closed under vector addition and scalar multiplication. That is, (1) for all $\mathbf{u}_1, \mathbf{u}_2 \in \mathcal{U}$, we also have $\mathbf{u}_1 + \mathbf{u}_2 \in \mathcal{U}$, and (2) for all $\mathbf{u} \in \mathcal{U}$ and $\alpha \in \mathbb{C}$, we have $\alpha\mathbf{u} \in \mathcal{U}$. (When only working with real numbers, replace $\alpha \in \mathbb{C}$ by $\alpha \in \mathbb{R}$.)

We now consider several important subspaces associated with a matrix $\mathbf{A} \in \mathbb{C}^{m\times n}$. The *range* (or *column space*) of $\mathbf{A}$ is defined as

$$\mathrm{Ran}(\mathbf{A}) = \{\mathbf{A}\mathbf{x} : \mathbf{x} \in \mathbb{C}^n\} \subseteq \mathbb{C}^m.$$

The *kernel* (or *null space*) of $\mathbf{A}$ is defined as

$$\mathrm{Ker}(\mathbf{A}) = \{\mathbf{x} \in \mathbb{C}^n : \mathbf{A}\mathbf{x} = \mathbf{0}\} \subseteq \mathbb{C}^n.$$

The *rank* of $\mathbf{A}$, denoted $\mathrm{rank}(\mathbf{A})$, is the dimension of the range of $\mathbf{A}$, i.e., the number of linearly independent vectors in a basis for $\mathrm{Ran}(\mathbf{A})$. We say the square matrix $\mathbf{A} \in \mathbb{C}^{n\times n}$ is *nonsingular* provided the following equivalent conditions hold:

- $\mathbf{A}^{-1}$ exists;
- $\mathrm{Ran}(\mathbf{A}) = \mathbb{C}^n$ (i.e., $\mathrm{rank}(\mathbf{A})) = n$);
- $\mathrm{Ker}(\mathbf{A}) = \{\mathbf{0}\}$;
- $\mathbf{A}$ has no zero eigenvalues.

The *identity matrix* is always denoted by $\mathbf{I}$, the *zero matrix* by $\mathbf{0}$.

The transpose of a vector $\mathbf{x} \in \mathbb{C}^n$ is denoted by $\mathbf{x}^T$, and the conjugate-transpose is denoted by $\mathbf{x}^*$; thus, $\mathbf{x}^* = \overline{\mathbf{x}}^T \in \mathbb{C}^{1\times n}$. Note that $\mathbf{x}^T$ and $\mathbf{x}^*$ are row vectors. (If $\mathbf{x}$ has only real entries, then $\mathbf{x}^T = \mathbf{x}^*$.) The conjugate-transpose generalizes to matrices, where for $\mathbf{A} \in \mathbb{C}^{m\times n}$, we have $\mathbf{A}^* = \overline{\mathbf{A}}^T \in \mathbb{C}^{n\times m}$.

Two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{C}^n$ are *orthogonal* provided $\mathbf{x}^*\mathbf{y} = 0$. Two subspaces $\mathcal{U} \subseteq \mathbb{C}^n$ and $\mathcal{V} \subseteq \mathbb{C}^n$ are orthogonal provided $\mathbf{u}^*\mathbf{v} = 0$ for all $\mathbf{u} \in \mathcal{U}$ and $\mathbf{v} \in \mathcal{V}$. If $\mathcal{U}$ and $\mathcal{V}$ are orthogonal, we write $\mathcal{U} \perp \mathcal{V}$. The set of all vectors $\mathbf{v} \in \mathbb{C}^n$ that are orthogonal to the subspace $\mathcal{U}$ is denoted by

$$\mathcal{U}^\perp = \{\mathbf{v} \in \mathbb{C}^n : \mathbf{v}^*\mathbf{u} = 0 \text{ for all } \mathbf{u} \in \mathcal{U}\}.$$

Suppose that $\mathbf{x} \in \mathrm{Ran}(\mathbf{A})$ and $\mathbf{z} \in \mathrm{Ker}(\mathbf{A}^*)$. Then there exists some $\mathbf{y} \in \mathbb{C}^n$ such that $\mathbf{x} = \mathbf{A}\mathbf{y}$, and we have $(\mathbf{z}^*\mathbf{x})^* = (\mathbf{z}^*\mathbf{A}\mathbf{y})^* = \mathbf{y}^*\mathbf{A}^*\mathbf{z} = \mathbf{y}^*\mathbf{0} = 0$, which implies $\mathrm{Ran}(\mathbf{A}) \perp \mathrm{Ker}(\mathbf{A}^*)$. Swapping the role of $\mathbf{A}$ and $\mathbf{A}^*$ gives $\mathrm{Ran}(\mathbf{A}^*) \perp \mathrm{Ker}(\mathbf{A})$. After considering the dimensions of these spaces, we arrive at what Gilbert Strang calls the *Fundamental Theorem of Linear Algebra*:

$$\begin{aligned} \mathbb{C}^m &= \mathrm{Ran}(\mathbf{A}) \oplus \mathrm{Ker}(\mathbf{A}^*), \qquad & \mathrm{Ran}(\mathbf{A}) \perp \mathrm{Ker}(\mathbf{A}^*) \\ \mathbb{C}^n &= \mathrm{Ran}(\mathbf{A}^*) \oplus \mathrm{Ker}(\mathbf{A}), \qquad & \mathrm{Ran}(\mathbf{A}^*) \perp \mathrm{Ker}(\mathbf{A}). \end{aligned}$$

The singular value decomposition, which will be covered in a few weeks, provides a natural means for untangling the four fundamental subspaces $\mathrm{Ran}(\mathbf{A})$, $\mathrm{Ran}(\mathbf{A}^*)$, $\mathrm{Ker}(\mathbf{A})$ and $\mathrm{Ker}(\mathbf{A}^*)$, and observing their orthogonality.

It is worth noting some important classes of square matrices:

- $\mathbf{A} \in \mathbb{C}^{n\times n}$ is *Hermitian* provided $\mathbf{A}^* = \mathbf{A}$. $\mathbf{A}$ is *symmetric* provided $\mathbf{A}^T = \mathbf{A}$. (For real matrices, these terms can be used interchangeably. However, some applications, e.g., in electromagnetics, give rise to complex symmetric matrices, which lack many of the fine properties enjoyed by Hermitian matrices.)
- $\mathbf{Q} \in \mathbb{C}^{n\times n}$ is *unitary* provided $\mathbf{Q}^*\mathbf{Q} = \mathbf{I}$. Since $\mathbf{Q}$ is square, we must have that $\mathbf{Q}^{-1} = \mathbf{Q}^*$, and hence also that $\mathbf{Q}\mathbf{Q}^* = \mathbf{I}$. A real unitary matrix is also called an *orthogonal matrix*. (Notice that if $\mathbf{Q}^*\mathbf{Q} = \mathbf{I} \in \mathbb{C}^{n\times n}$ for the rectangular matrix $\mathbf{Q} \in \mathbb{C}^{m\times n}$ with $m > n$, then $\mathbf{Q}\mathbf{Q}^* \neq \mathbf{I} \in \mathbb{C}^{m\times m}$. Can you explain why?)

1.1.1. Vector and matrix norms.

As we study numerical *analysis*, we shall require a means of measuring distance.

Definition. A function $\|\cdot\| : \mathbb{C}^n \to \mathbb{R}$ is a *norm* provided:

- $\|\mathbf{x}\| \geq 0$ for all $\mathbf{x} \in \mathbb{C}^n$; $\|\mathbf{x}\| = 0$ if and only if $\mathbf{x} = \mathbf{0}$ (positivity);
- $\|\alpha\mathbf{x}\| = |\alpha|\|\mathbf{x}\|$ for all $\alpha \in \mathbb{C}$ and $\mathbf{x} \in \mathbb{C}^n$ (scaling);
- $\|\mathbf{x} + \mathbf{y}\| \leq \|\mathbf{x}\| + \|\mathbf{y}\|$ for all $\mathbf{x}, \mathbf{y} \in \mathbb{C}^n$ (triangle inequality).

In this course, the only vector norms we shall use are the *p-norms*:

$$\|\mathbf{x}\|_p = \Big(\sum_{j=1}^n |x_j|^p\Big)^{1/p}$$

for $p \ge 1$. The values $p = 1, 2, \infty$ are by far the most common:

$$\|\mathbf{x}\|_1 = \sum_{j=1}^n |x_j|; \qquad \|\mathbf{x}\|_2 = \Big(\sum_{j=1}^n |x_j|^2\Big)^{1/2} = \sqrt{\mathbf{x}^*\mathbf{x}}; \qquad \|\mathbf{x}\|_\infty = \max_{j=1,\ldots,n} |x_j|.$$

Among innumerable useful norm relationships, pride of place belongs to (1) the *Cauchy–Schwarz* inequality:

$$|\mathbf{x}^*\mathbf{y}| \;\le\; \|\mathbf{x}\|_2\|\mathbf{y}\|_2,$$

with equality holding when $\mathbf{x}$ and $\mathbf{y}$ are collinear, and (2) the *Pythagorean Theorem*: for orthogonal $\mathbf{x}$ and $\mathbf{y}$,

$$\|\mathbf{x}+\mathbf{y}\|_2^2 = \|\mathbf{x}\|_2^2 + \|\mathbf{y}\|_2^2.$$

The vector 2-norm enjoys another important property: it is *unitarily invariant*: for any unitary matrix $\mathbf{U} \in \mathbb{C}^{n\times n}$,

$$\|\mathbf{U}\mathbf{x}\|_2 = \|\mathbf{x}\|_2$$

for all $\mathbf{x} \in \mathbb{C}^n$. This is easy to prove: since $\mathbf{U}^*\mathbf{U} = \mathbf{I}$,

$$\|\mathbf{U}\mathbf{x}\|_2^2 = (\mathbf{U}\mathbf{x})^*(\mathbf{U}\mathbf{x}) = \mathbf{x}^*\mathbf{U}^*\mathbf{U}\mathbf{x} = \mathbf{x}^*\mathbf{x} = \|\mathbf{x}\|_2^2.$$

This fact has a natural 'physical' interpretation: If $\mathbf{U} = [\mathbf{u}_1\ \mathbf{u}_2\ \cdots\ \mathbf{u}_n]$ is unitary, then its columns $\mathbf{u}_1, \ldots, \mathbf{u}_n$ form an orthonormal basis for $\mathbb{C}^n$. The vector $\mathbf{U}\mathbf{x} = \sum_{j=1}^n x_j\mathbf{u}_j$ is a representation of the vector $\mathbf{x}$ in the coordinate system whose axes are given by $\mathbf{u}_1, \ldots, \mathbf{u}_n$. The statement $\|\mathbf{U}\mathbf{x}\|_2 = \|\mathbf{x}\|_2$ simply means, 'the length of $\mathbf{x}$ does not change when we convert from the standard orthonormal basis (columns of the identity matrix) to the new orthonormal basis (columns of $\mathbf{U}$).'

We can use norms to measure the magnitude of a matrix. The same axioms stated above for vector norms apply here.

Definition. A function $\|\cdot\| : \mathbb{C}^{m\times n} \to \mathbb{R}$ is a *matrix norm* provided:

- $\|\mathbf{A}\| \ge 0$ for all $\mathbf{A} \in \mathbb{C}^{m\times n}$; $\|\mathbf{A}\| = 0$ if and only if $\mathbf{A} = \mathbf{0}$ (positivity);
- $\|\alpha\mathbf{A}\| = |\alpha|\|\mathbf{A}\|$ for all $\alpha \in \mathbb{C}$ and $\mathbf{A} \in \mathbb{C}^{m\times n}$ (scaling);
- $\|\mathbf{A}+\mathbf{B}\| \le \|\mathbf{A}\| + \|\mathbf{B}\|$ for all $\mathbf{A}, \mathbf{B} \in \mathbb{C}^{m\times n}$ (triangle inequality).

The most important class of norms are the *induced* matrix norms, which are defined in terms of some vector norm $\|\cdot\|$:

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

That is, induced matrix norms measure the maximum amount a matrix can stretch a vector beyond its original length. It is often handy to use the equivalent definition:

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

The matrix norms induced by the vector p-norms are particularly useful. When $p = 1$ and $p = \infty$, we have the simple formulas

$$\|\mathbf{A}\|_1 = \max_{1\le k\le n} \sum_{j=1}^{m} |a_{jk}|; \qquad \|\mathbf{A}\|_\infty = \max_{1\le j\le m} \sum_{k=1}^{n} |a_{jk}|.$$

That is, the 1-norm is the maximum absolute column sum, while the ∞-norm is the maximum absolute row sum. Most useful for many applications is the induced matrix 2-norm (called by some the *spectral norm*):

$$\|\mathbf{A}\|_2 = \max_{\mathbf{x}\neq\mathbf{0}} \frac{\|\mathbf{A}\mathbf{x}\|_2}{\|\mathbf{x}\|_2}.$$

The matrix 2-norm inherits *unitary invariance* from the vector 2-norm: for any unitary matrices $\mathbf{U}$ and $\mathbf{V}$, $\|\mathbf{U}\mathbf{A}\mathbf{V}\|_2 = \|\mathbf{A}\|_2$. We shall derive a formula for the 2-norm from the singular value decomposition in a few weeks:

$$\|\mathbf{A}\|_2 = \max\{\sqrt{\lambda} : \lambda \text{ is an eigenvalue of } \mathbf{A}^*\mathbf{A}\}.$$

For values of p other than 1, 2, and ∞, there is no simple formula for the induced matrix p-norm.

Trefethen and Bau present a very nice graphical description of induced matrix norms in their Figure 3.1 (page 20), which we highly recommend. (See `norm_demo.m` on the class website.)

The useful *Frobenius norm* is not induced by any vector norm:

$$\|\mathbf{A}\|_F = \Big(\sum_{j=1}^{m}\sum_{k=1}^{n} |a_{jk}|^2\Big)^{1/2}.$$

Many matrix norms are *submultiplicative*,

$$\|\mathbf{A}\mathbf{B}\| \le \|\mathbf{A}\|\,\|\mathbf{B}\|,$$

a property that enables much analysis. (In fact, some authors add this condition as a fourth requirement in the definition of a 'matrix norm'.) To prove that the induced matrix norms are submultiplicative, observe that

$$\|\mathbf{A}\mathbf{B}\| = \max_{\mathbf{x}\neq\mathbf{0}} \frac{\|\mathbf{A}\mathbf{B}\mathbf{x}\|}{\|\mathbf{x}\|} = \max_{\mathbf{x}\neq\mathbf{0}} \frac{\|\mathbf{A}\mathbf{B}\mathbf{x}\|}{\|\mathbf{B}\mathbf{x}\|}\frac{\|\mathbf{B}\mathbf{x}\|}{\|\mathbf{x}\|} \le \Big(\max_{\mathbf{x}\neq\mathbf{0}} \frac{\|\mathbf{A}\mathbf{B}\mathbf{x}\|}{\|\mathbf{B}\mathbf{x}\|}\Big)\Big(\max_{\mathbf{x}\neq\mathbf{0}} \frac{\|\mathbf{B}\mathbf{x}\|}{\|\mathbf{x}\|}\Big) \le \|\mathbf{A}\|\|\mathbf{B}\|.$$

(Can you explain both inequalities, and confirm that division by zero is not a major concern?)

The Frobenius norm, too, is submultiplicative. However, there exist norms that satisfy the three basic matrix norm axioms, but are not submultiplicative, e.g.,

$$\|\mathbf{A}\| = \max_{j,k} |a_{jk}|$$

satisfies the positivity, scaling, and triangle inequality properties, yet one can construct $\mathbf{A}$ and $\mathbf{B}$ such that $\|\mathbf{A}\mathbf{B}\| > \|\mathbf{A}\|\,\|\mathbf{B}\|$. (Try it!)

► An extensive discussion of vector and matrix norms can be found in Chapter 5 of *Matrix Analysis* by R. A. Horn and C. R. Johnson, Cambridge, 1985. For a sophisticated treatment of the class of *unitarily invariant norms*, see Chapter 4 of *Matrix Analysis* by R. Bhatia, Springer, 1997.